

Recent Progress in Computational Science at NERSC

Horst D. Simon

**Director, NERSC Center and Computational Research
Lawrence Berkeley National Laboratory**



San Diego, CA

September 4, 2003



Berkeley Lab



Founded in 1931 by E.O. Lawrence on the Berkeley Campus;
Moved to Current Site in 1940



Outline



- **NERSC and Computational Research at Berkeley Labs**
- **Brief Overview of NERSC Systems and Users**
- **Projects**
 - Data management
 - Grids
 - Architectures
- **Computational Science Results**





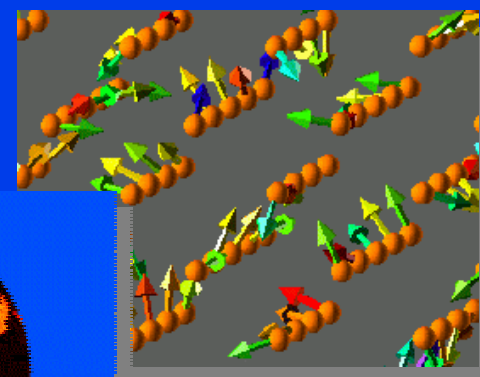
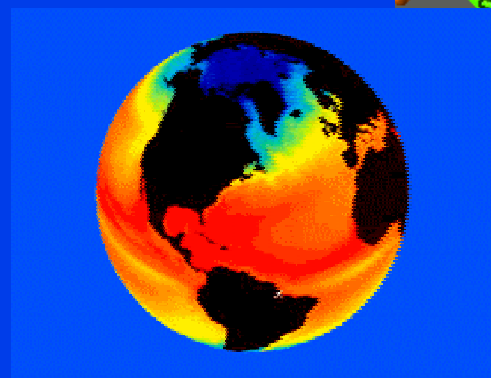
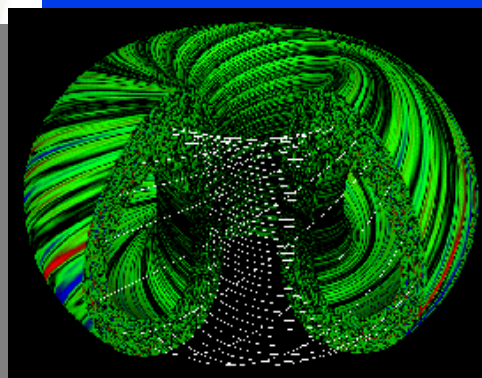
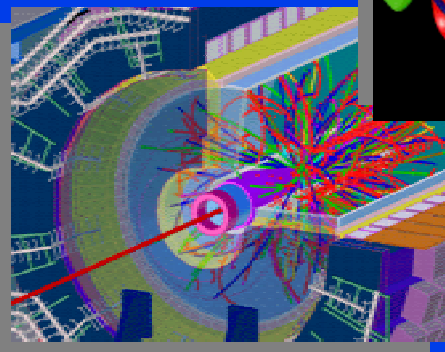
NERSC Center Division

Horst D. Simon, Division Director

**William T.C. Kramer, Division Deputy
and Facility General Manager**

National Energy Research Scientific Computing Center

- Serves all disciplines of the DOE Office of Science
- ~2000 Users in ~400 projects
- Focus on large-scale computing





NERSC Center Overview



- Funded by DOE, annual budget \$28M, about 65 staff
- Supports open, unclassified, basic research
- Located in the hills next to University of California, Berkeley campus
- close collaborations between university and NERSC in computer science and computational science
- close collaboration with about 125 scientists in the Computational Research Division at LBNL





Computational Research Division

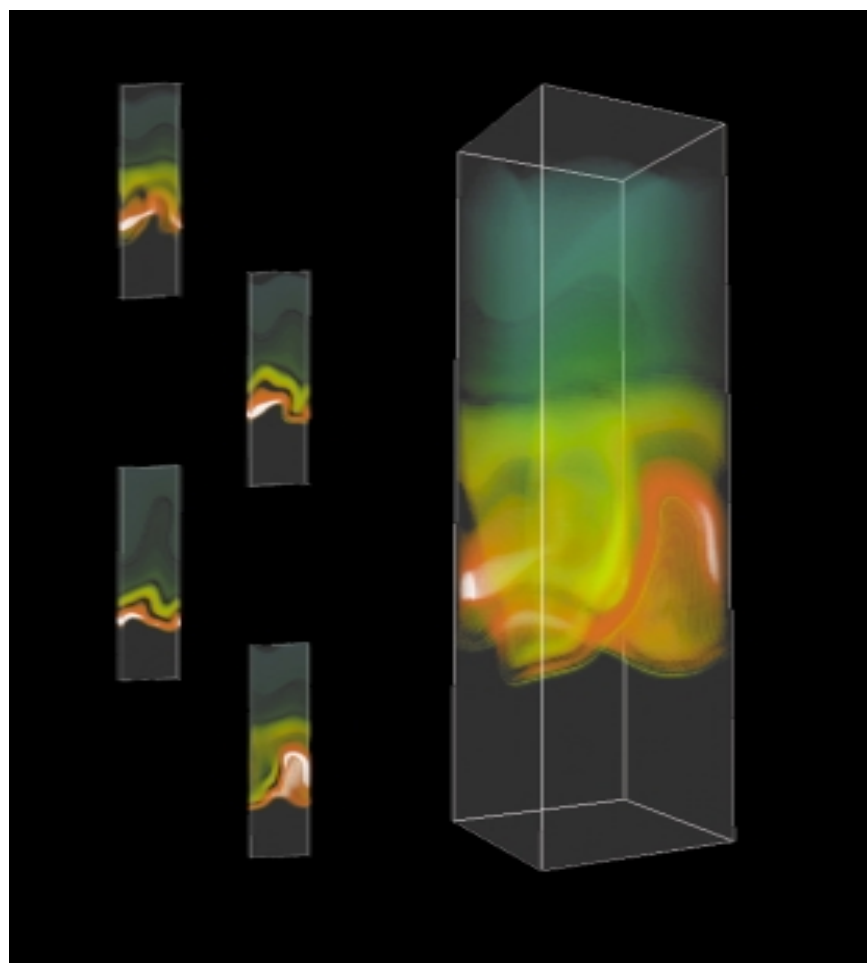
Horst D. Simon, Division Director

**Juan C. Meza, HPC Research Dept.
Head**

**William Johnston, Distributed Systems
Dept. Head**



Computational Research Division



Horst Simon, Division Director

The Computational Research Division(CRD) creates computational tools and techniques that enable scientific breakthroughs, by conducting applied research and development in computer science, computational science, and applied mathematics. CRD consists of two departments:

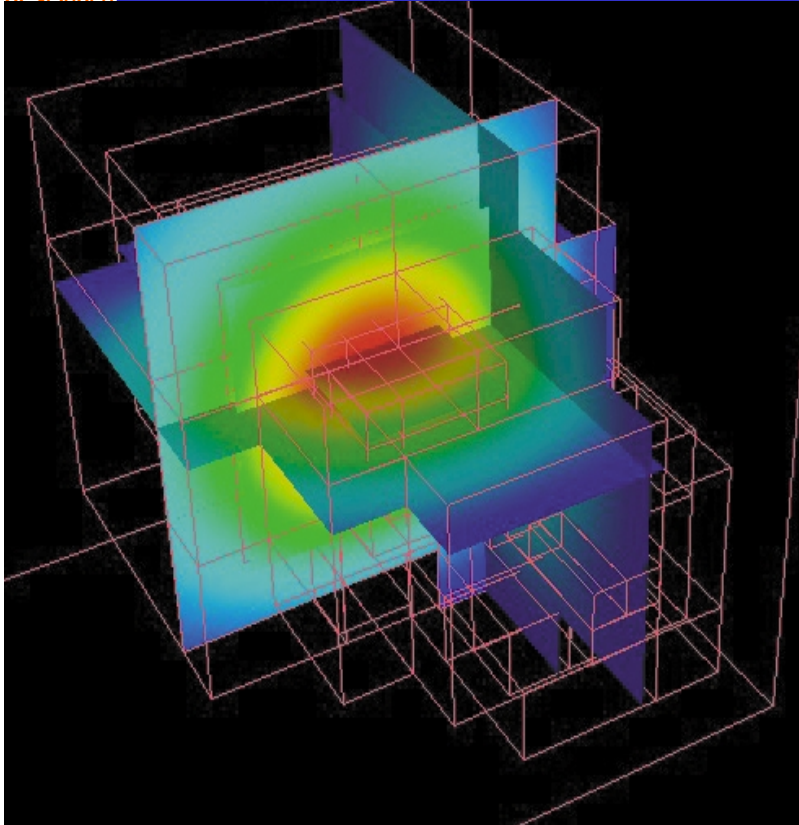
1. High Performance Computing Research
2. Distributed Systems

• About \$5M of SciDAC funded projects





High Performance Computing Research Department (HPCRD)



Juan Meza, Department Head

Groups:

- Applied Numerical Algorithms
- Center for Computational Sciences and Engineering
- Future Technologies
- Imaging and Informatics
- Scientific Computing
- Scientific Data Management
- Visualization

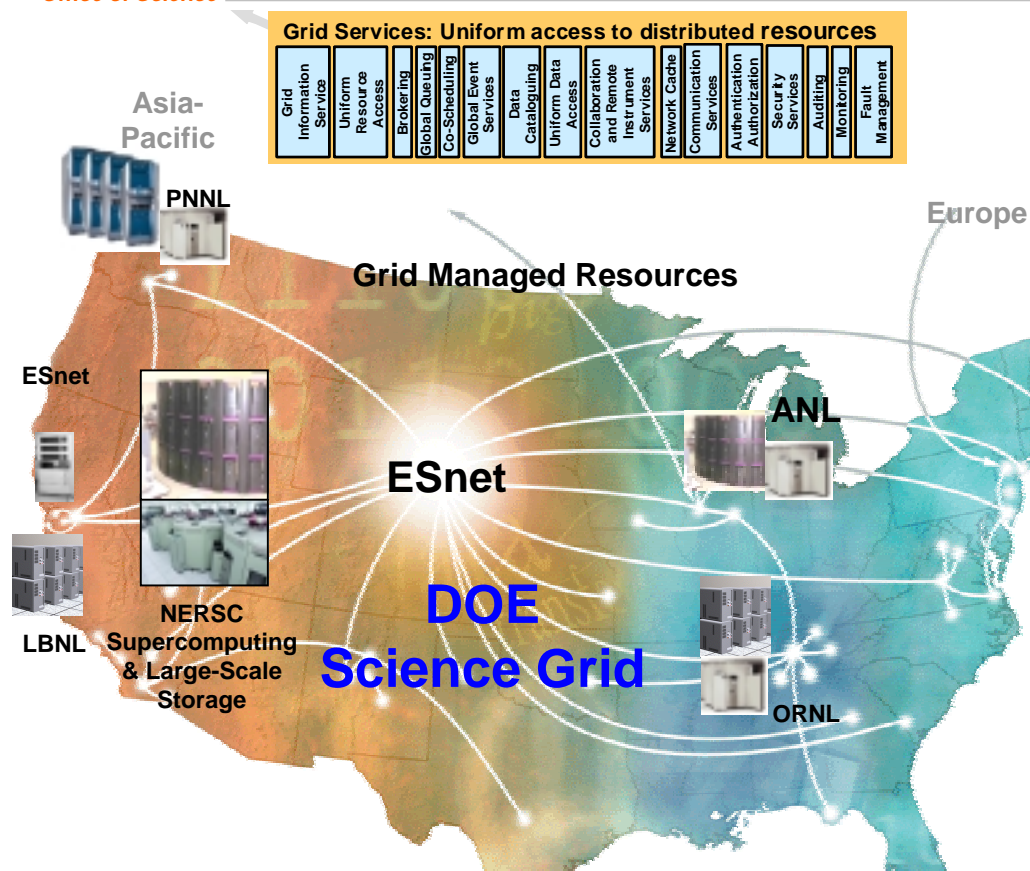
Total Staff: 108

The High Performance Computing Research Department conducts research and development in mathematical modeling, algorithmic design, software implementation, and system architectures, and evaluates new and promising technologies.





Distributed Systems Department



William Johnston, Department Head

Deb Agarwal, Department Deputy

Groups:

- Collaboration Technologies
- Data Intensive Distributed Computing
- Network Technologies
- Secure Grid Technologies

Total Staff: 25

The Distributed Systems Department researches and develops software components that allow scientists to address complex and large-scale computing and data analysis problems in a distributed environment such as the DOE Science Grid.





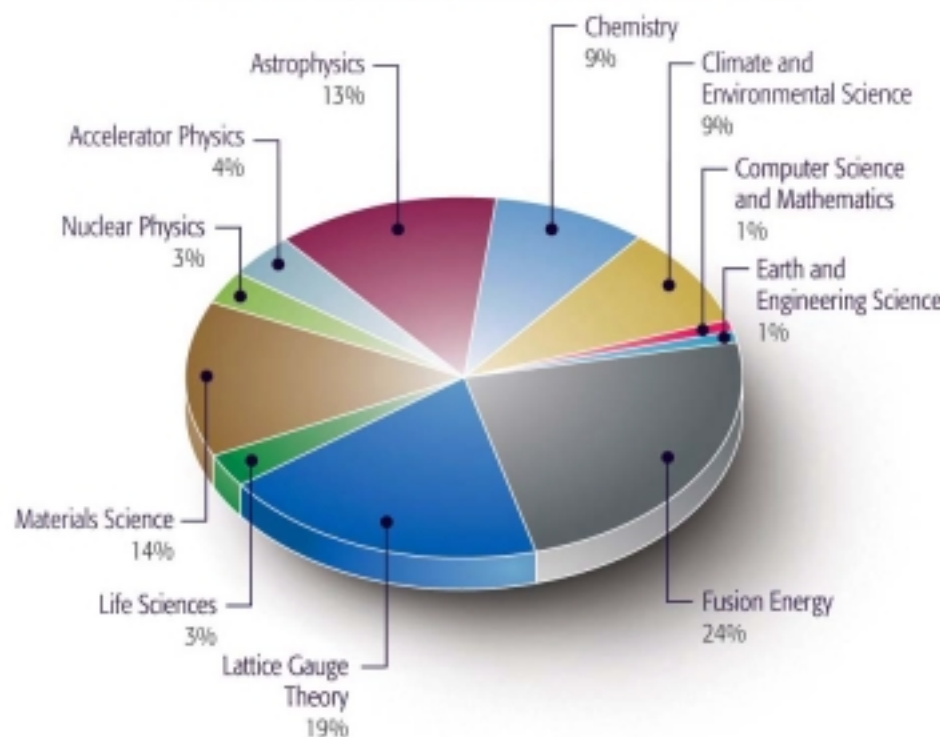
NERSC Users and Allocations



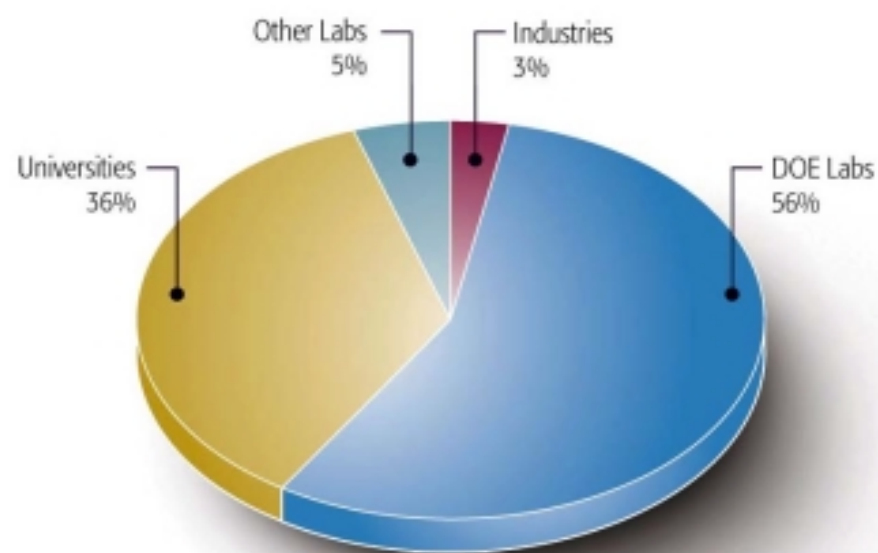
Usage by Discipline and Institution Type



NERSC Usage by Scientific Discipline, FY02



NERSC Usage by Institution Type, FY02

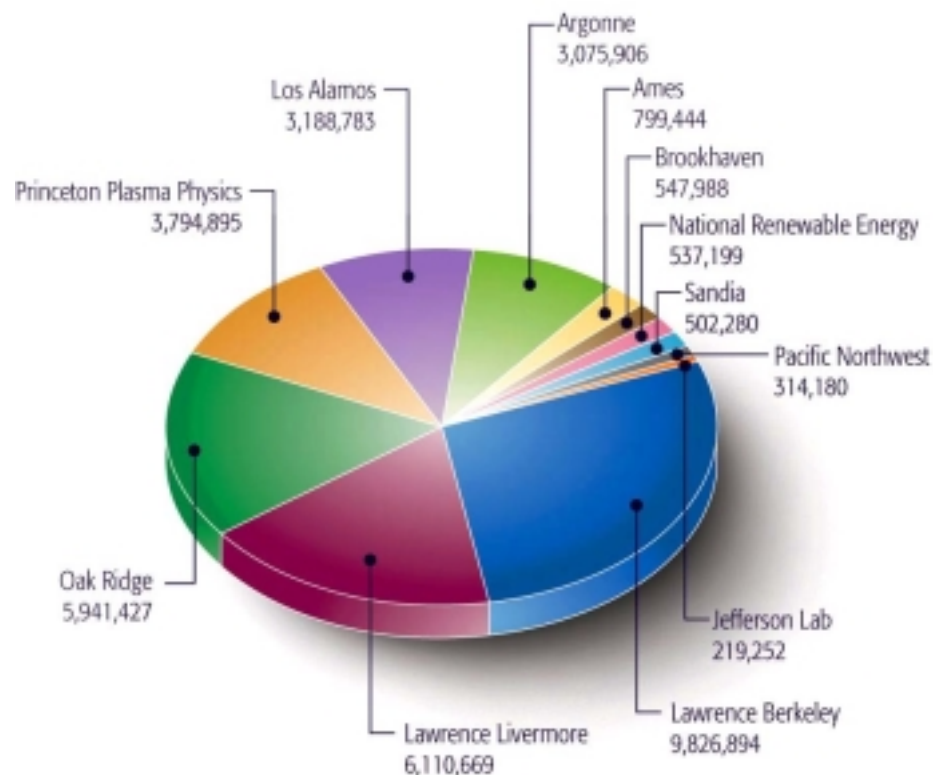




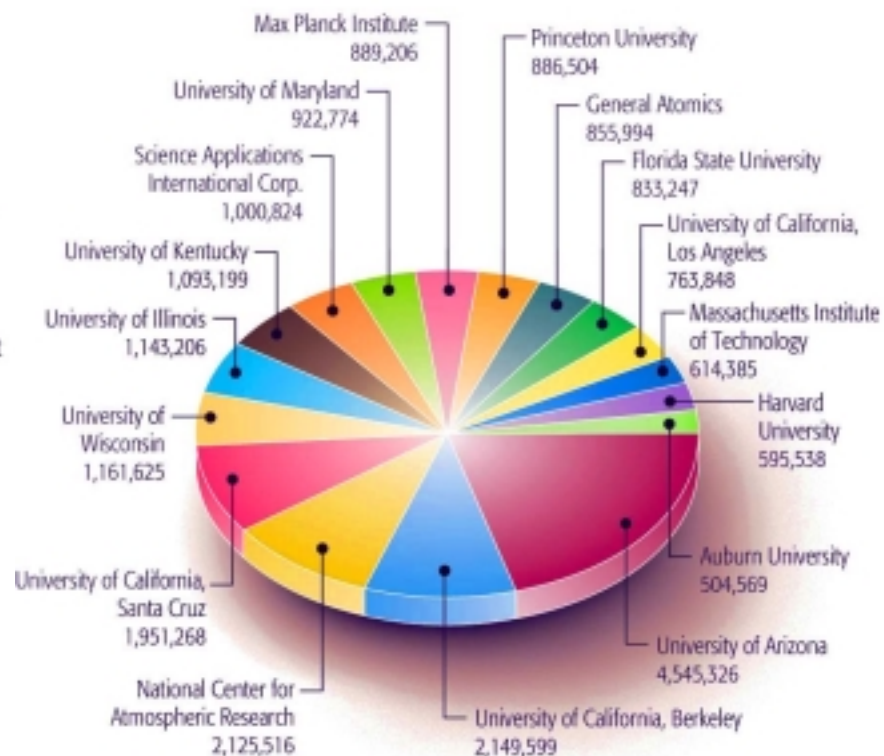
Usage by Institution



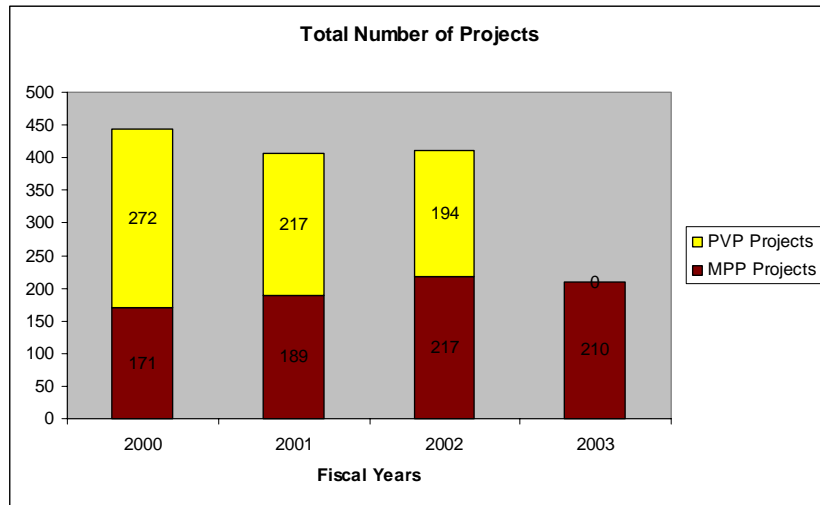
**Leading DOE Laboratory Usage at NERSC, FY02
(>200,000 processor hours)**



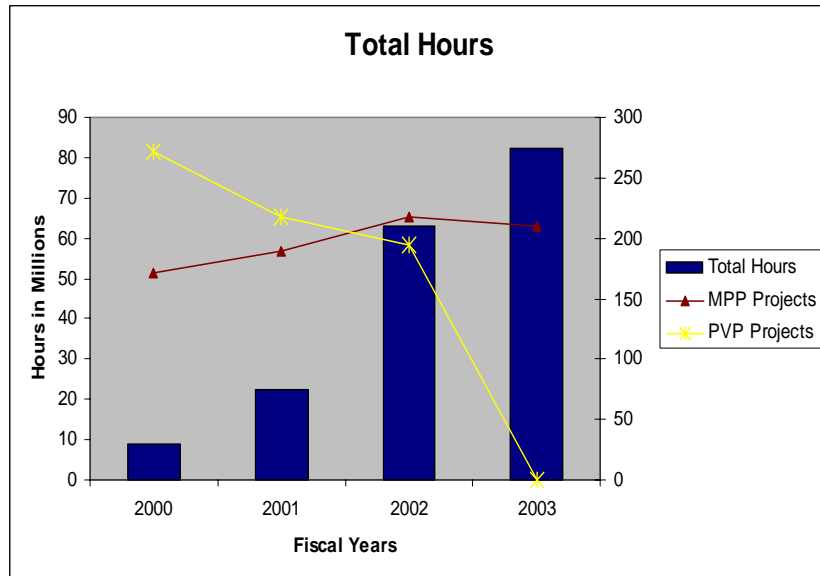
**Leading Academic and Related Usage at NERSC, FY02
(>500,000 processor hours)**



Increase in Capability Computing



The number of projects at NERSC has significantly decreased.



The amount of available hours has significantly increased



NERSC Systems



Upgraded NERSC 3E Characteristics



The upgraded NERSC 3E system has

- 416 16-way Power 3+ nodes with each CPU at 1.5 Gflop/s
 - 380 for computation
- 6,656 CPUs – 6,080 for computation
- Total Peak Performance of 10 Teraflop/s
- Total Aggregate Memory is 7.8 TB
- Total GPFS disk will be 44 TB
 - Local system disk is an additional 15 TB
- Combined SSP-2 is greater than 1.238 Tflop/s
- NERSC 3E is in full production as of March 1, 2003
 - nodes arrived in the first two weeks of November
 - Acceptance end of December 2002
 - 30-day availability test near completed Feb. 2003





TOP500 List of Most Powerful Computers



Rank	Manufacturer	Computer	R_{\max} [TF/s]	Installation Site	Country	Year	Area of Installation	# Proc
1	NEC	Earth-Simulator	35.86	Earth Simulator Center	Japan	2002	Research	5120
2	HP	ASCI Q, AlphaServer SC	13.88	Los Alamos National Laboratory	USA	2002	Research	8192
3	Linux Networx/ Quadrics	MCR Cluster	7.63	Lawrence Livermore National Laboratory	USA	2002	Research	2304
4	IBM	ASCI White SP Power3	7.3	Lawrence Livermore National Laboratory	USA	2000	Research	8192
5	IBM	Seaborg SP Power 3	7.3	NERSC Lawrence Berkeley Nat. Lab.	USA	2002	Research	6656
6	IBM/Quadrics	XSeries Cluster Xeon 2.4 GHz	6.59	Lawrence Livermore National Laboratory	USA	2003	Research	1920
7	Fujitsu	PRIMEPOWER HPC2500	5.41	National Aerospace Laboratory of Japan	Japan	2002	Research	2304
8	HP	rx2600 Itanium2 Cluster Quadrics	4.88	Pacific Northwest National Laboratory	USA	2003	Research	1536
9	HP	AlphaServer SC ES45 1 GHz	4.46	Pittsburgh Supercomputing Center	USA	2001	Academic	3016
10	HP	AlphaServer SC ES45 1 GHz	3.98	Commissariat a l'Energie Atomique (CEA)	France	2001	Research	2560



A Short Update

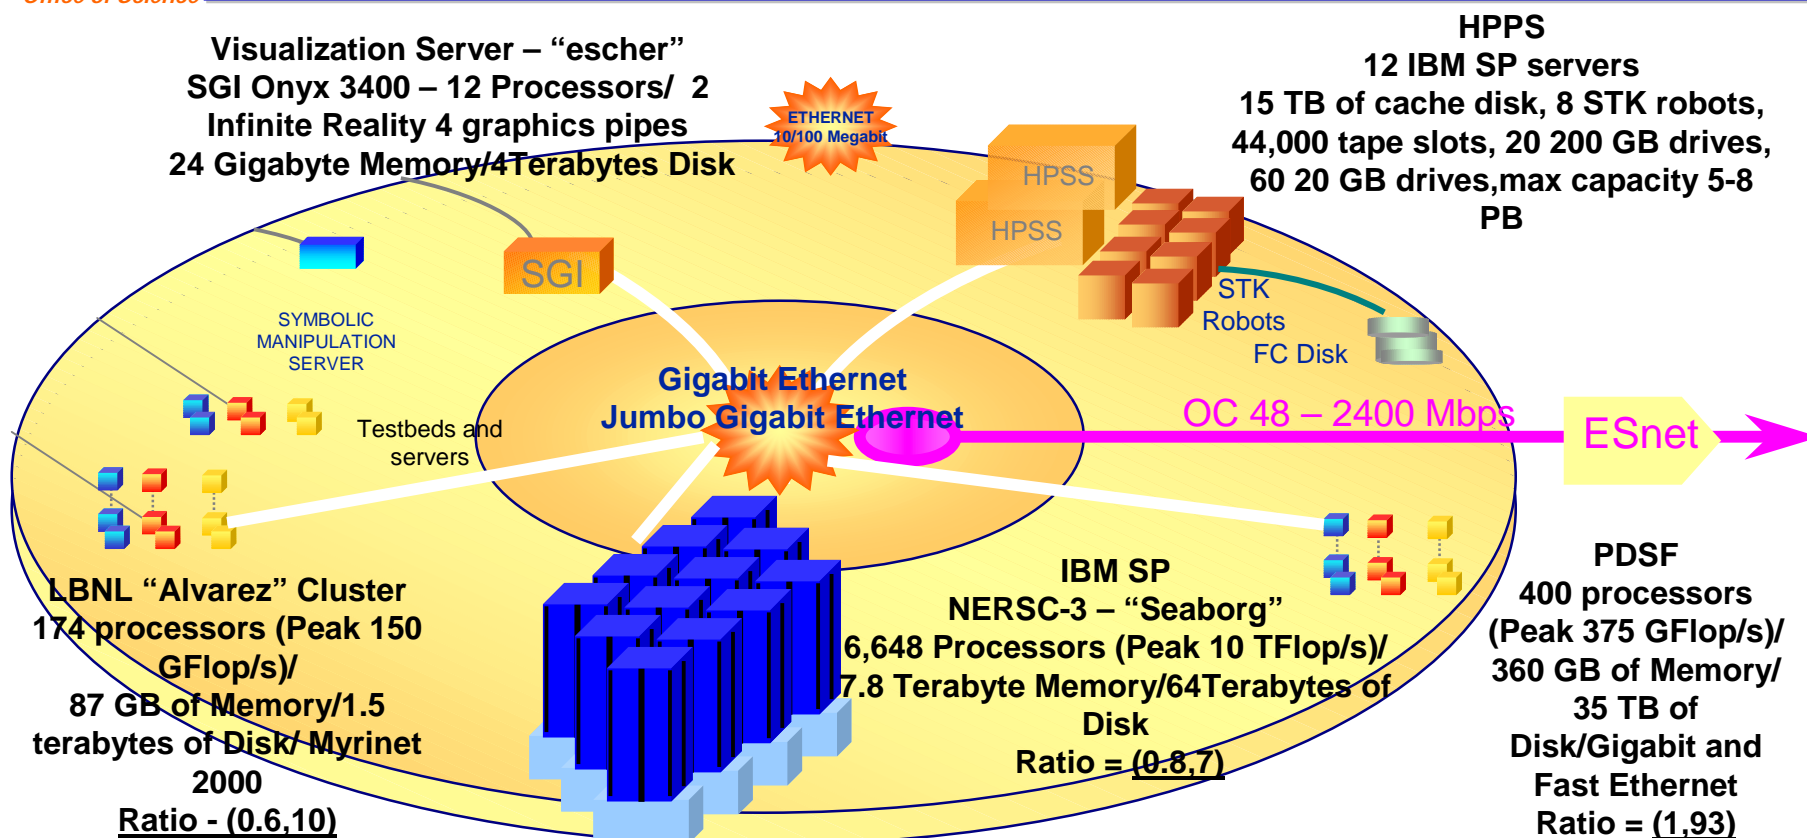


- **Expanded the IBM SP - Seaborg – in an unconventional decision**
 - In production one month earlier than planned
 - Added 60% more MPP hours to allocations in FY 03
 - Doubles allocations in FY 04 due to full year of operation
- **Scalability program and reimbursement program**
 - Readjusted queues on Seaborg to emphasize large jobs
- **Implemented OC-48 link to ESnet**
- **Added 200 GB tape drives to HPSS to bring it to 9 petabyte capacity**
 - Done to manage tape media costs (increased 400+ percent)





NERSC System Architecture

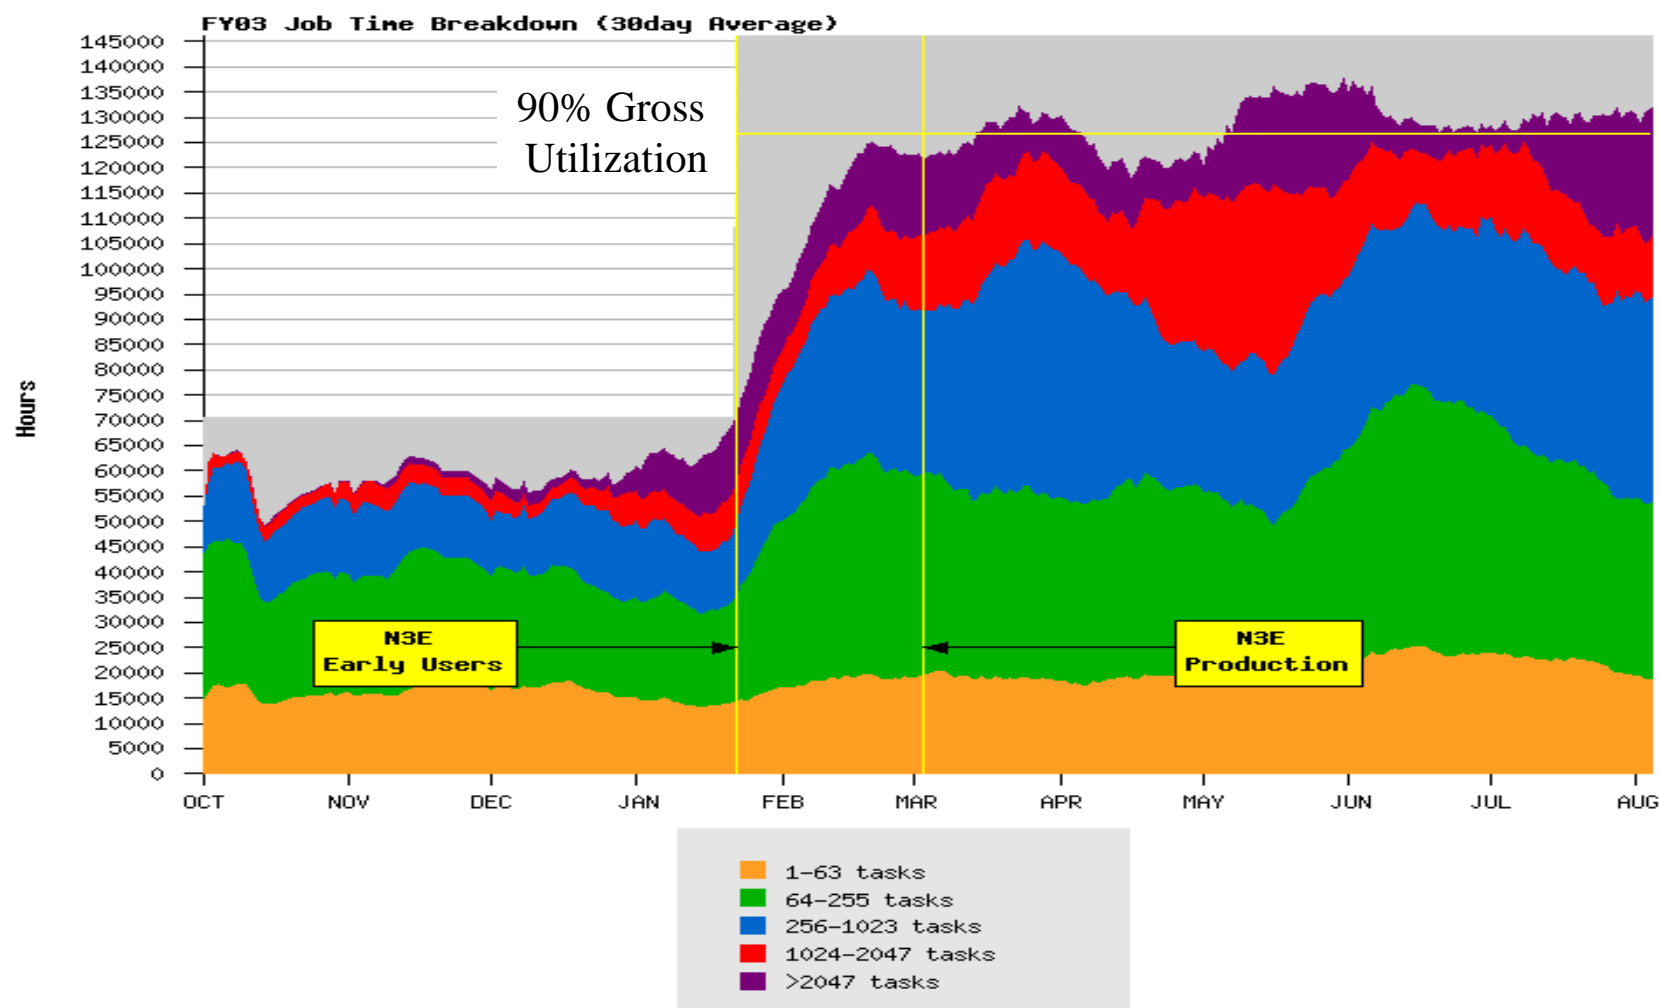


Ratio = (RAM Bytes per Flop, Disk Bytes per Flop)



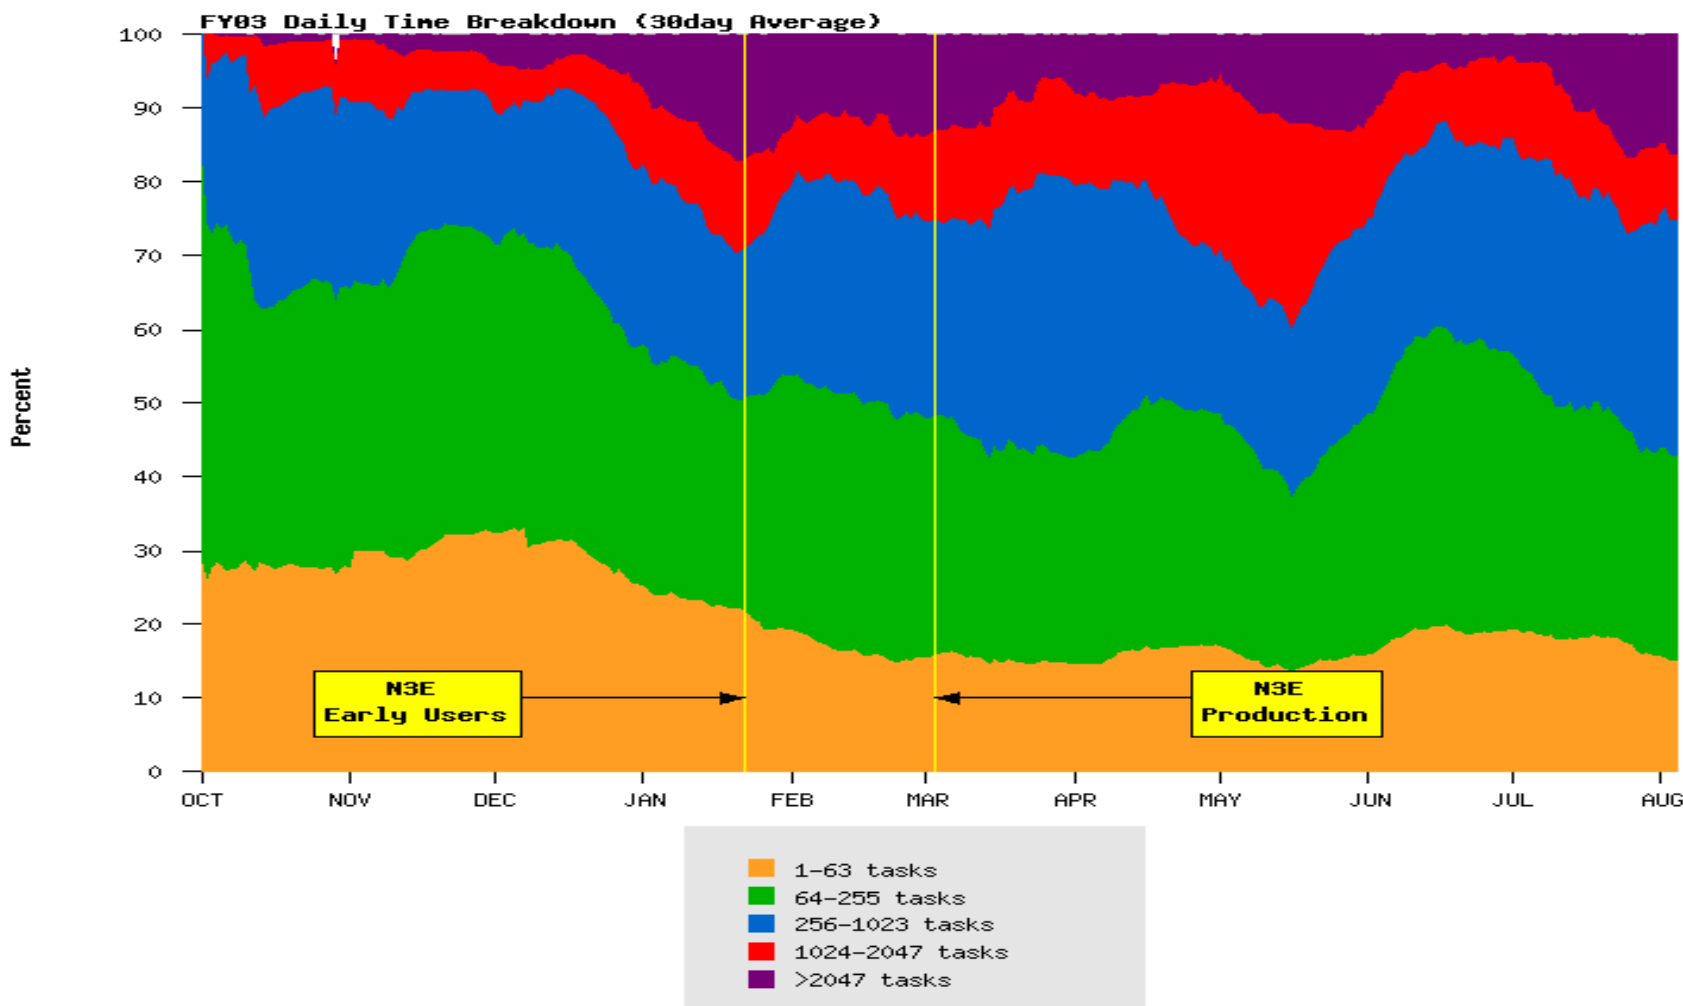


Large Job Sizes Run Regularly



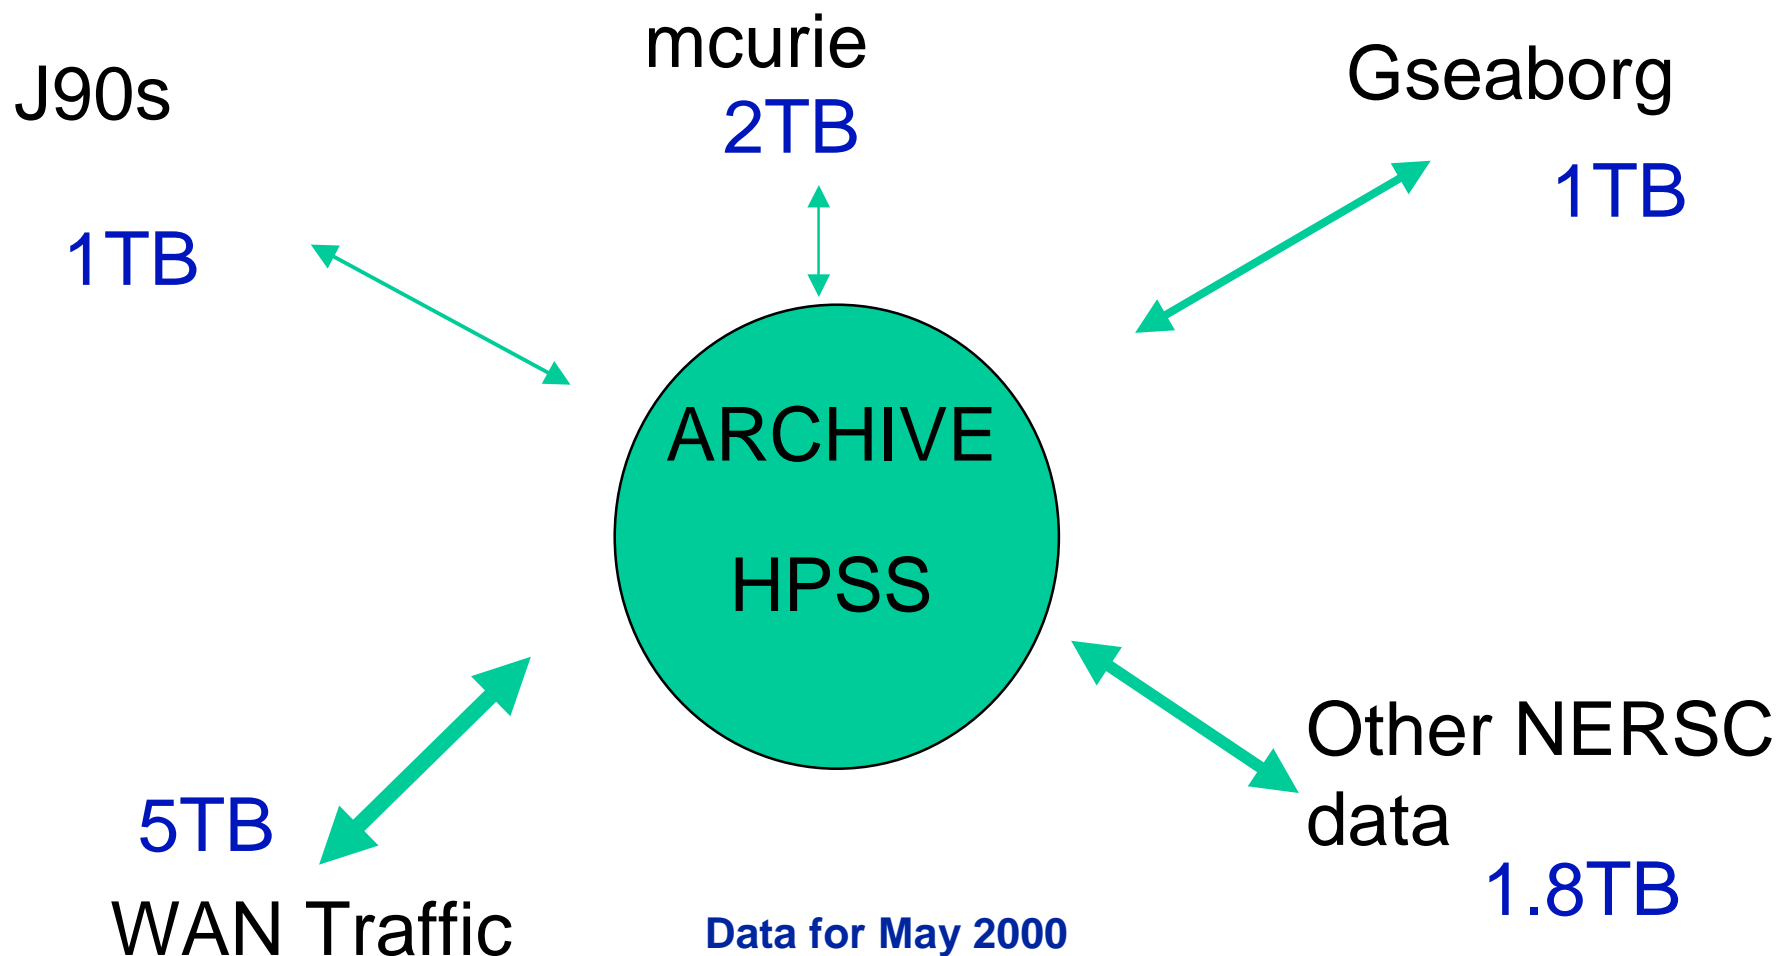


40% of the System used by >511 CPU Jobs (for 2003)



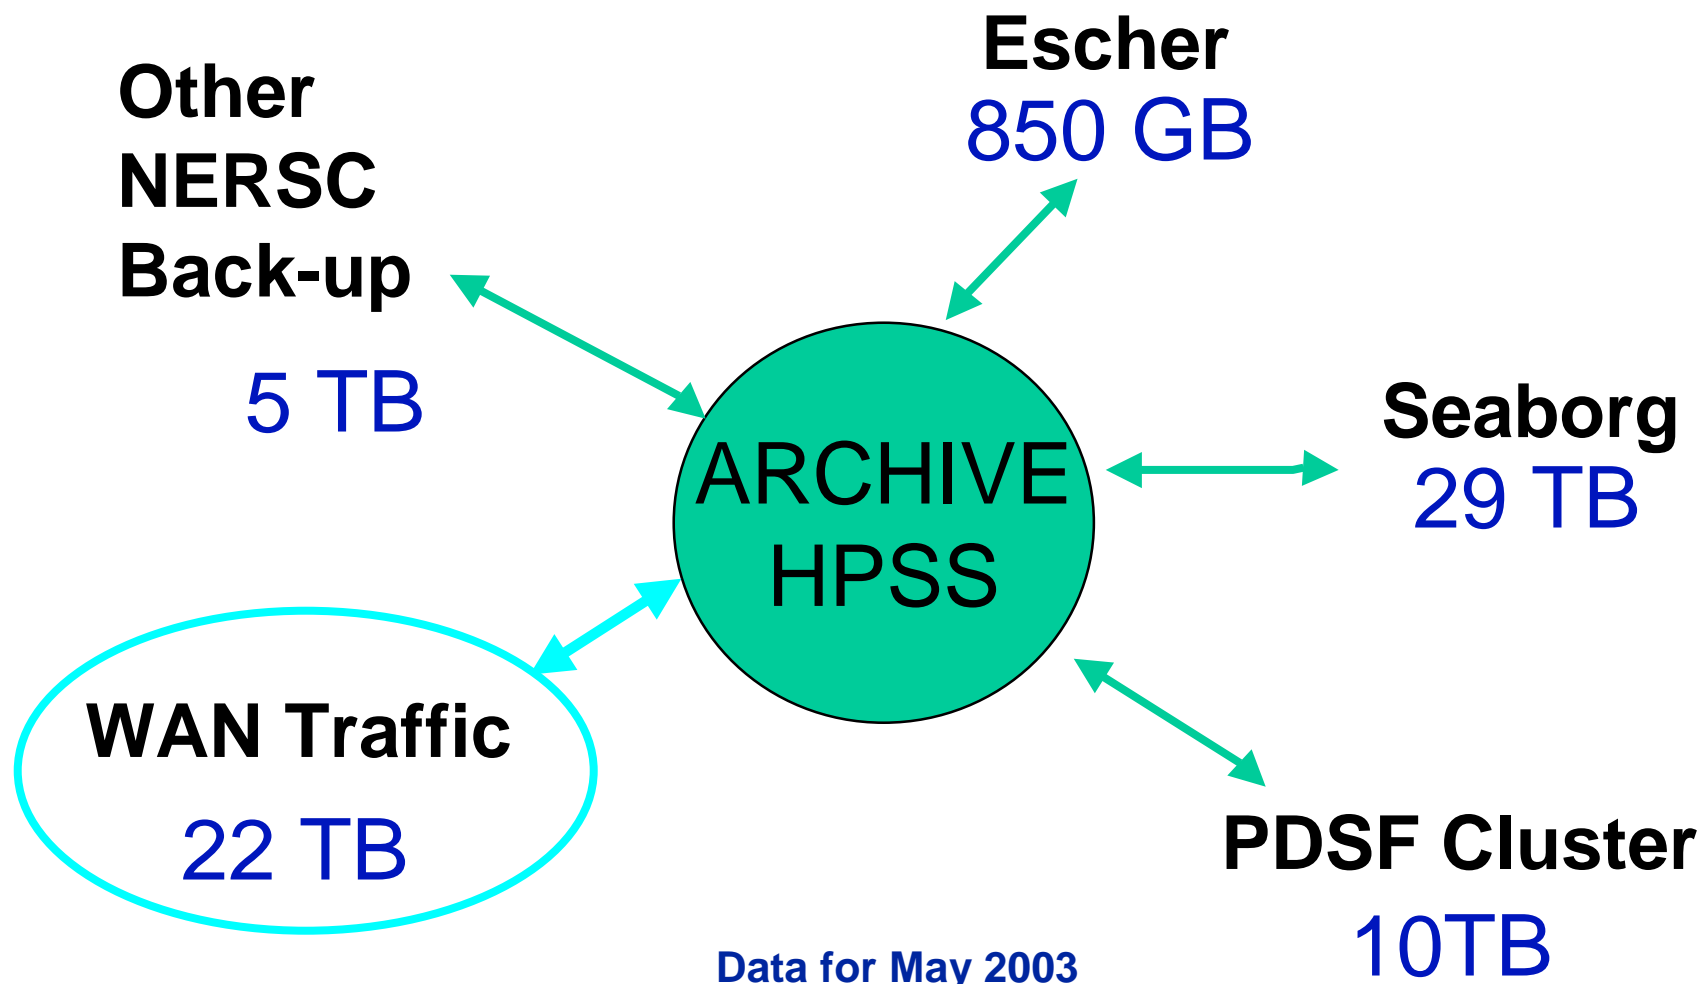
HPSS I/O Activity

May 2000 – Total = 10.8 TB



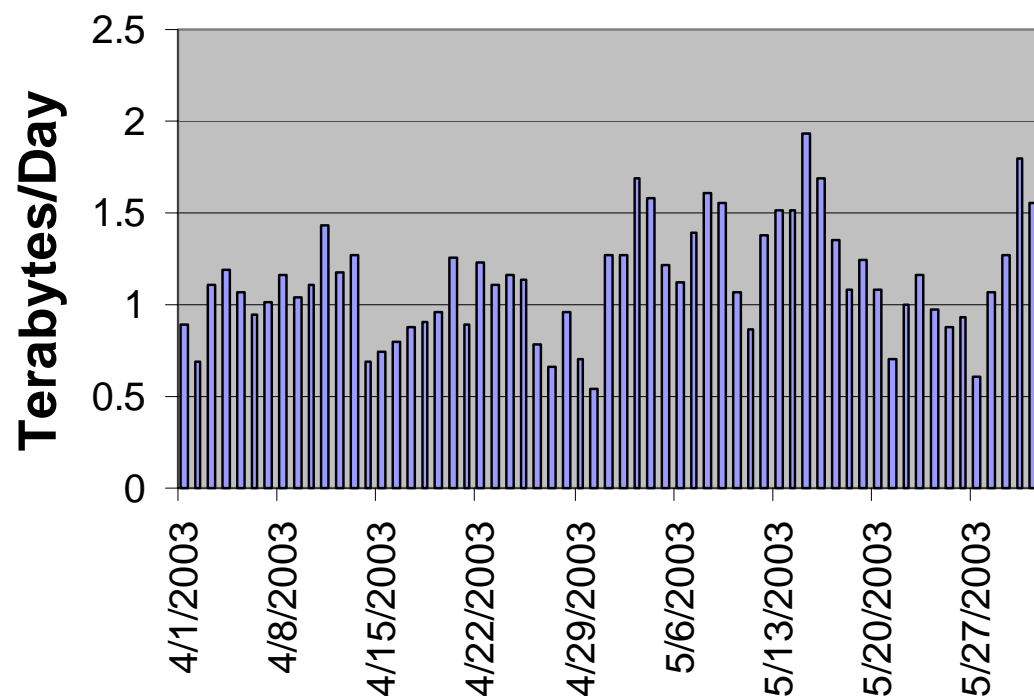


Monthly I/O Activity to Storage by Platform = 57 TB



NERSC is a Net Importer of Traffic

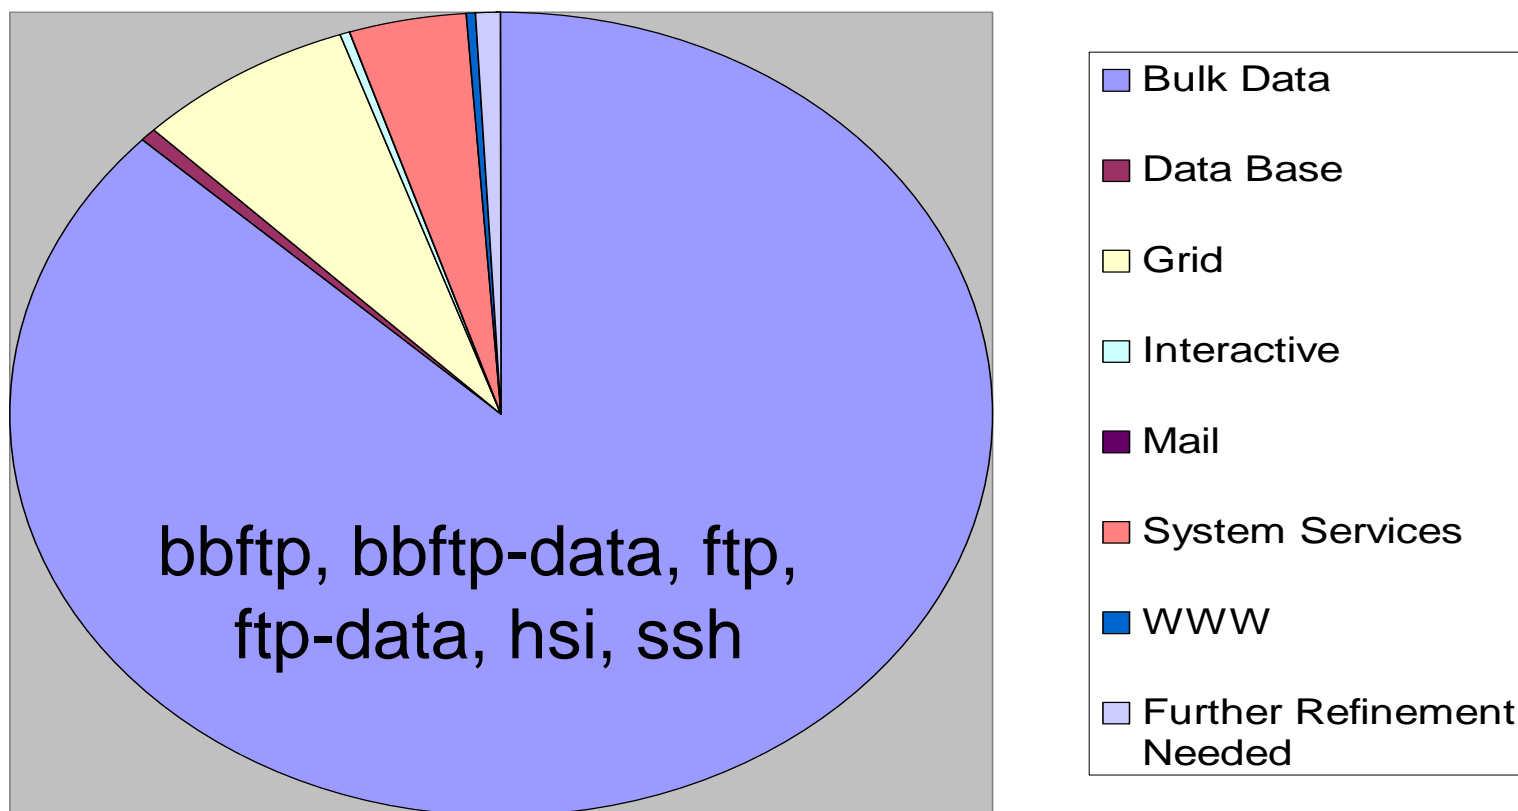
NERSC Border Traffic



- Traffic across the NERSC border:
 - April 2003 - 29.5 TB
 - May 2003 - 39.4 TB
- NERSC traffic accounts for approximately 20% of total ESNet traffic
- 76% of the NERSC traffic is inbound



Majority of NERSC Data Transfer is Bulk Data





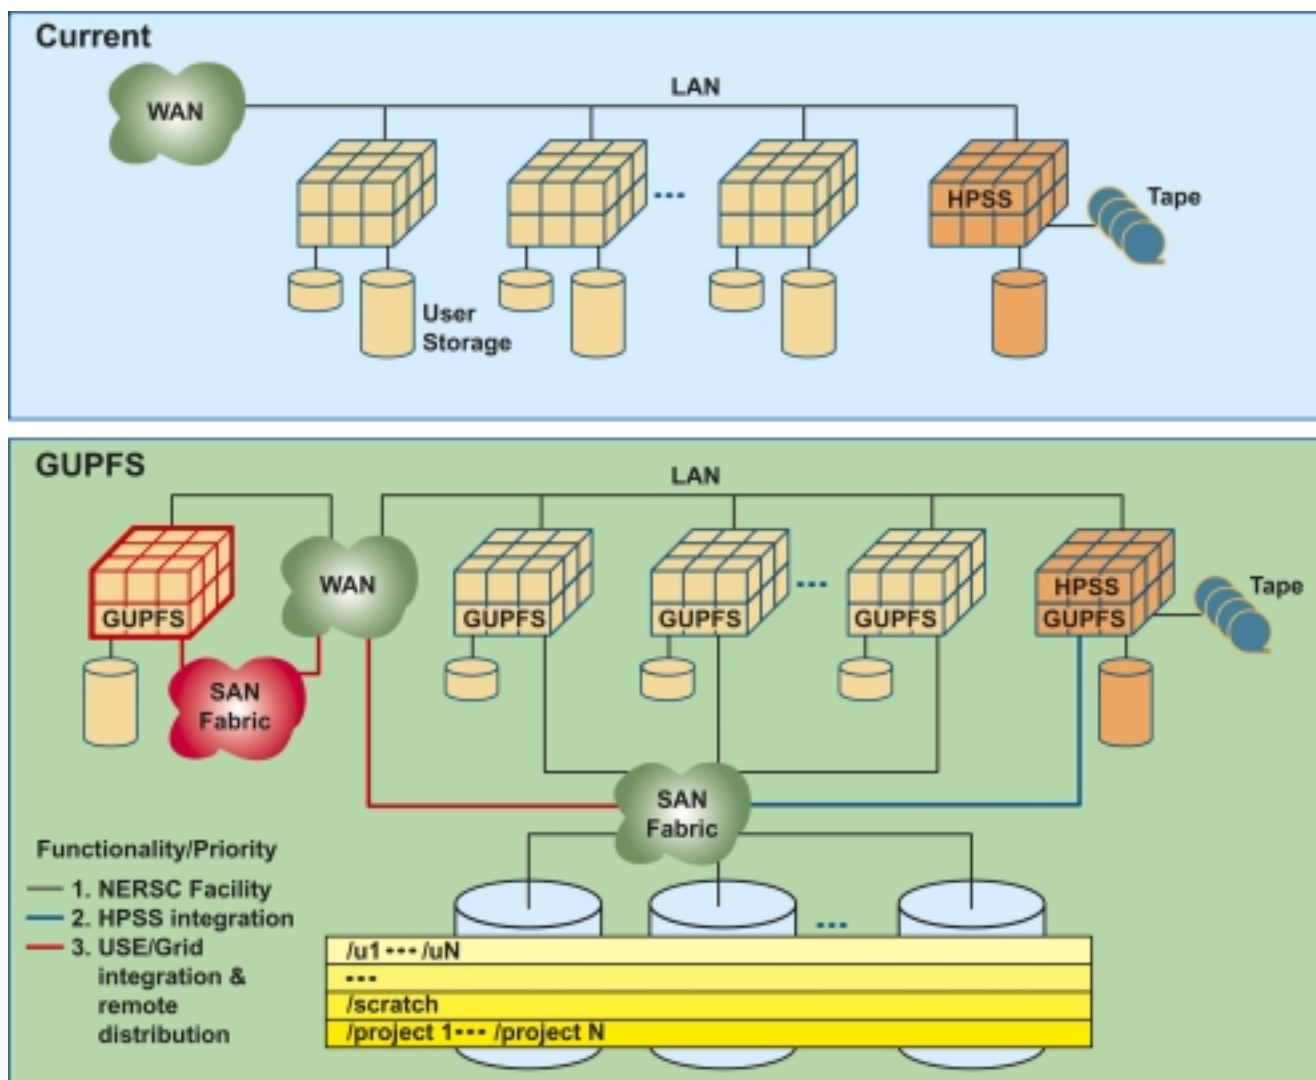
NERSC Projects



Global Unified Parallel File System



Concept Diagram



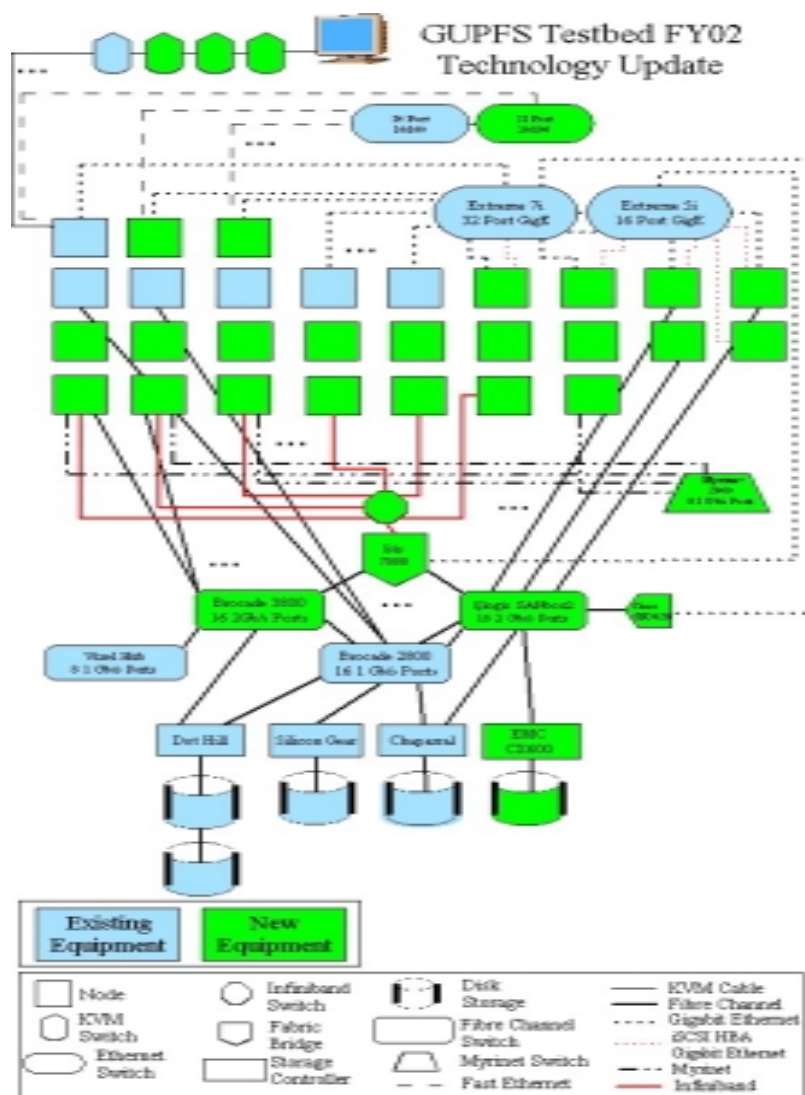


Global Parallel Universal File System



- 2002
 - Technology evaluation testbed.
 - Evaluate different shared-disk filesystems and SAN technologies (e.g., InfiniBand, Yotta-Yotta).
 - Establish preliminary evaluation criteria and develop shared-disk filesystem specific benchmarks accordingly.
 - Conduct preliminary wide-area network (WAN) storage distribution tests.
- 2003
 - Begin HPSS integration.
 - Continue evaluation of shared-disk filesystem and SAN technologies (e.g., ASCI Path Forward SGSFS, other filesystems, iSCSI).
 - Implement second generation of SAN fabrics and technology evaluation testbed.
 - Complete identification of evaluation criteria.
 - Start scalability tests on auxiliary systems.
- 2004
 - Begin HPSS HSM integration.
 - Conduct ASCI Path Forward SGSFS beta testing.
 - Select GUPFS solution shortlist candidates.
 - Begin final evaluation and selection.
- 2005
 - Complete final evaluation and select GUPFS solution.
 - Generate deployment plans for consolidated storage and GUPFS.
 - Begin phased consolidated storage buildup.
 - Begin phased GUPFS deployment.
- 2006
 - Continue phased storage buildup.
 - Continue phased GUPFS deployment.
 - GUPFS production on NERSC production system.
- Post 2006 - GUPFS is a standard architectural component of NERSC.





- Working with Cray, Inc and Unlimited Scale as well



GUPFS Recent Activities



- **Built a flexible testbed**
 - Storage devices, interconnect fabric and filesystem can be tested and developed in a modular fashion
- **Developed testing methodology and two flexible parallel file system benchmarks**
 - Test I/O Performance, Metadata performance, reliability and emulation of applications
 - MPTIO and METABENCH benchmarks
- **Evaluated storage devices**
 - Yotta Yotta
 - 3PARdata
 - DataDirect Network
 - EMC
 - DotHill
 - SiliconGear
- **Evaluated interconnect fabrics**
 - **Fibre Channel**
 - QLogic FC HBA
 - QLogic FC switch
 - Brocade FC switch
 - Cisco iSCSI/FC switch
 - **iSCSI**
 - Cisco iSCSI switch
 - Intel iSCSI HBA
 - Adaptec iSCSI HBA and TOE card
 - QLogic iSCSI HBA
 - **InfiniBand**
 - InfiniCon switch and HCA
 - Topspin switch and HCA
- **First initial evaluation of file system software**
 - Sestina gfs
 - IBM GPFS
 - Lustre
 - Maximum Throughput InfiniARRAY File System
 - Ibrix File System

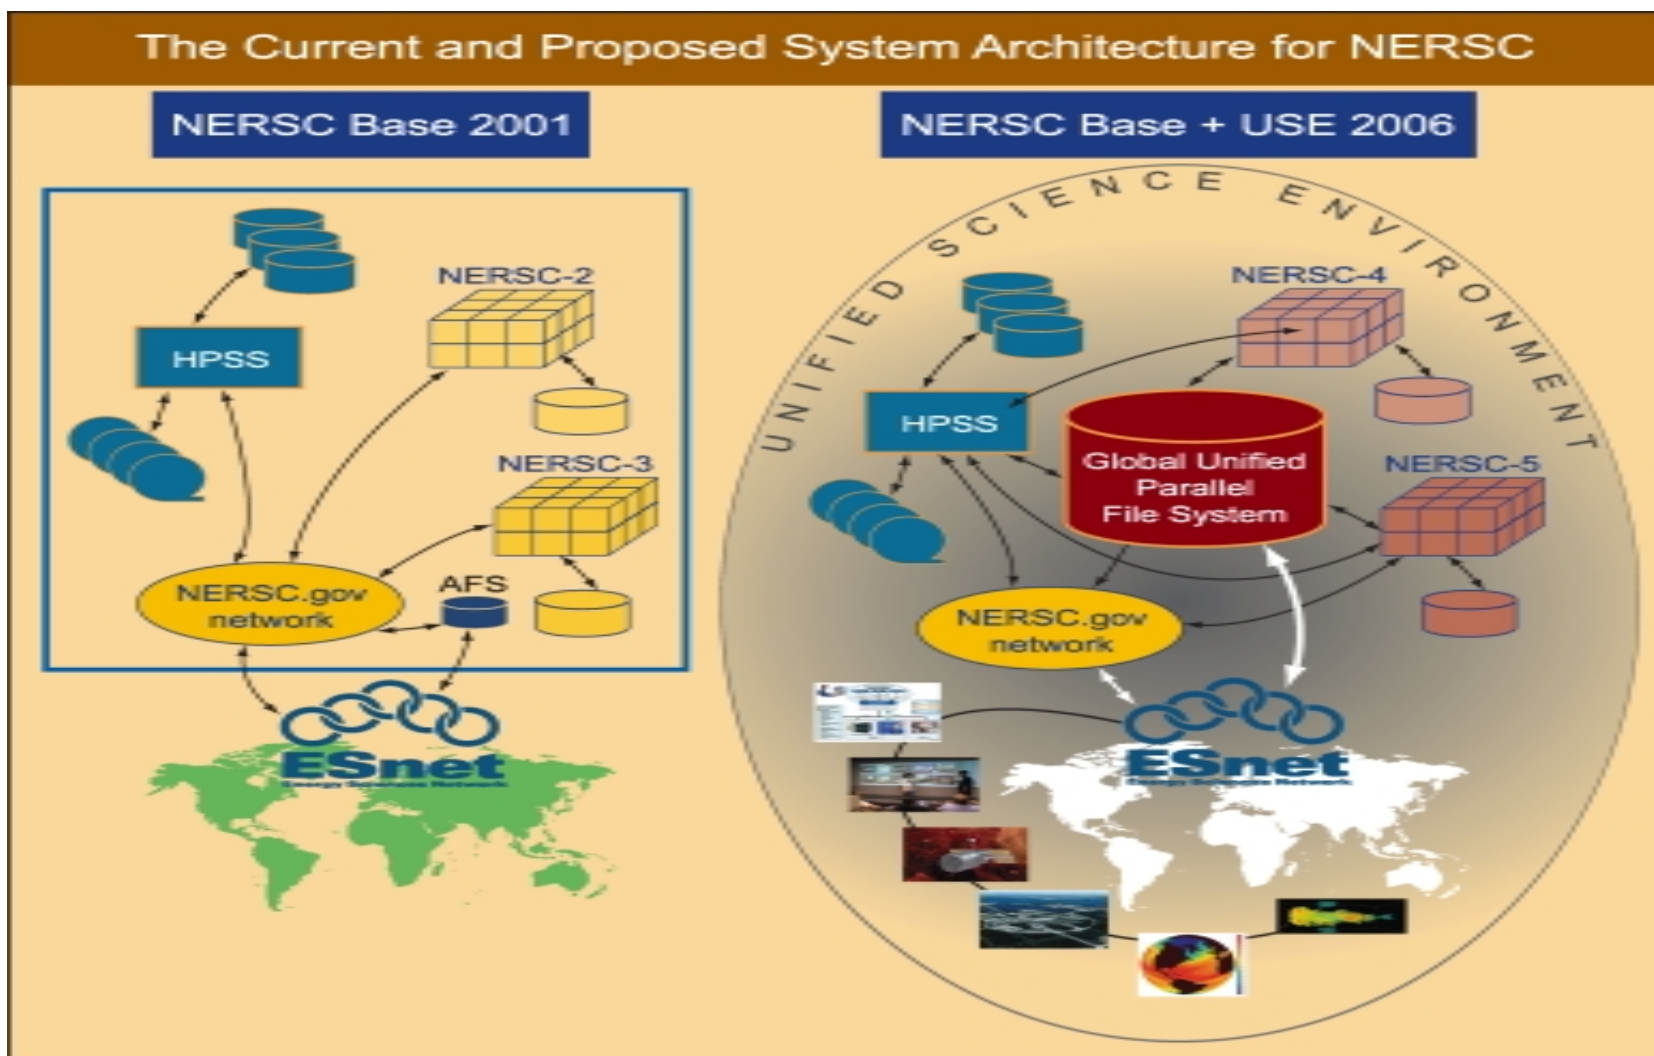




Connecting to DOE Science Grid



NERSC Systems Will Evolve





NERSC and the Grid



- **Multi-year plan**
 - 2002
 - Data Grid pre-production activities
 - Track computational grid, collaboration, and workflow development
 - Established a collaborative agreement with IBM to accelerate deploying Grid Technology
 - 2003
 - Focus on data Grid production rollout
 - Pre-production compute Grid
 - Track collaboration and workflow development
 - Earth Systems Grid Prototype
 - 2004
 - Focus on compute Grid production rollout
 - Pre-production collaboration and workflow
 - FY2005
 - Focus on collaboration and workflow production rollout
 - FY2006
 - All major [USE/GRID](#) components on NERSC production systems





Latest Activities



- **Infrastructure is in place on all system**
 - LDAP, CAs, basic globus functionality, etc.
 - Working in cooperation with IBM to test, improve and field GTK 2.2 on the IBM SP – now a few early beta users have access
 - Testing the Grid with firewalls
 - Implement Grid aware IDS features
- **Production use of the Grid for Storage and PDSF**
- **Developed an interim solution to grid enable HPSS.**
 - Now being distributed until the new GridFTP for HPSS is available
- **White papers**
 - Security
 - Implementation Issues
 - Vision for HPSS and the Grid





NERSC and Blue Planet

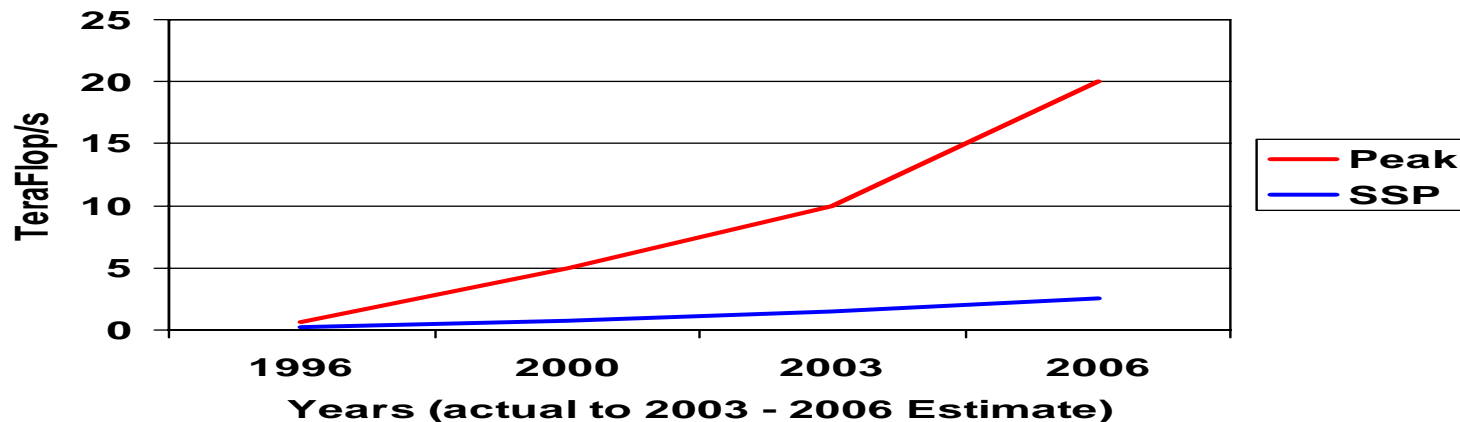


The Divergence Problem



- The requirements of high performance computing for science and engineering and the requirements of the commercial market are diverging.
- The commercial-clusters-of-COTS-SMPs approach is no longer sufficient to provide the highest level of performance
 - Lack of memory bandwidth
 - High interconnect latency
 - Lack of interconnect bandwidth
 - Lack of high performance parallel I/O
 - High cost of ownership for large scale systems

Divergence





What is Blue Planet



- ***"Blue Planet"*** is a "science driven" design process to develop systems that are simultaneously more effective for science and sustainable and cost effective for vendors.
 - White Paper uses IBM as an example of what can be done with this process
 - Can be applied to a number of vendors
- **Blue Planet** is a new concept for a sustainable computer architecture more effective for science and engineering applications
 - A specific implementation leveraging the IBM roadmap that better balances scientific processing needs and the commercial viability
 - Described as a "ultrascale" scale system on the order of the Earth Simulator

<http://www.nersc.gov/news/blueplanet.html>

and

<http://www.nersc.gov/news/ArchDevProposal.5.01.pdf>





Approach



- Engage the vendor community with a new approach to leveraging their major R&D/product roadmaps to create new architectures that are much more effective for science
 - Study applications critical to DOE Office of Science and others. For example:
 - Material Science, Combustion simulation and adaptive methods, Computational astrophysics, Nanoscience (new drugs and also new microchip technologies), Biochemical and Biosciences (protein folding/interactions), Climate modeling, High Energy Physics (particle accelerators and astrophysics), Multi-grid Eigen solvers and LA methods
 - Identify key bottlenecks found in these critical applications
 - Outline a high level approach to address the challenges
 - Follow-up meetings for detailed drill down by the vendor experts, computer scientists and application scientists at NERSC
 - Iterate on proposed solution





Needs Based on Scientific Applications



	AMR	Coupled Climate	Astrophysics		Nanoscience	
			MADCAP	Cactus	FLAPW	LSMS
Sensitive to global bisection	X	X	X		X	
Sensitive to processor to memory latency	X	X			X	
Sensitive to network latency	X	X	X	X	X	
Sensitive to point to point communications	X	X				X
Sensitive to OS interference in frequent barriers				X	X	
Benefits from deep CPU pipelining	X	X	X	X	X	X
Benefits from Large SMP nodes	X					





Full IBM Blue Planet System Components



- New IH++ Wide Node - 8 CPUs per node
 - POWER5 GS Processor - 2.5GHz
 - Single core MCM
- 2048 node system (8 Nodes per frame)
 - 16K processors @ 10GF per CPU = 160TF Peak
- Virtual Vector Architecture - VIVA
- Federation Switch - 3 stage topology
 - 8GB/s per server for the uni direction communication bandwidth.
- 40-50 TF Sustained on 2-3 selected applications
- 256 TB of memory = 16GB per CPU
 - May reduced to 128TB of memory if it can sustain full memory BW
- 2.5PB disk in I/O system [approximately 48 IO nodes]
- Approximately 600 Frames
 - 256 compute racks, 250 Disk racks, 160 Switch racks
 - 12,000-15,000 Sq Feet; 5-7 MWatts Power
- Scientists will focus on application optimization

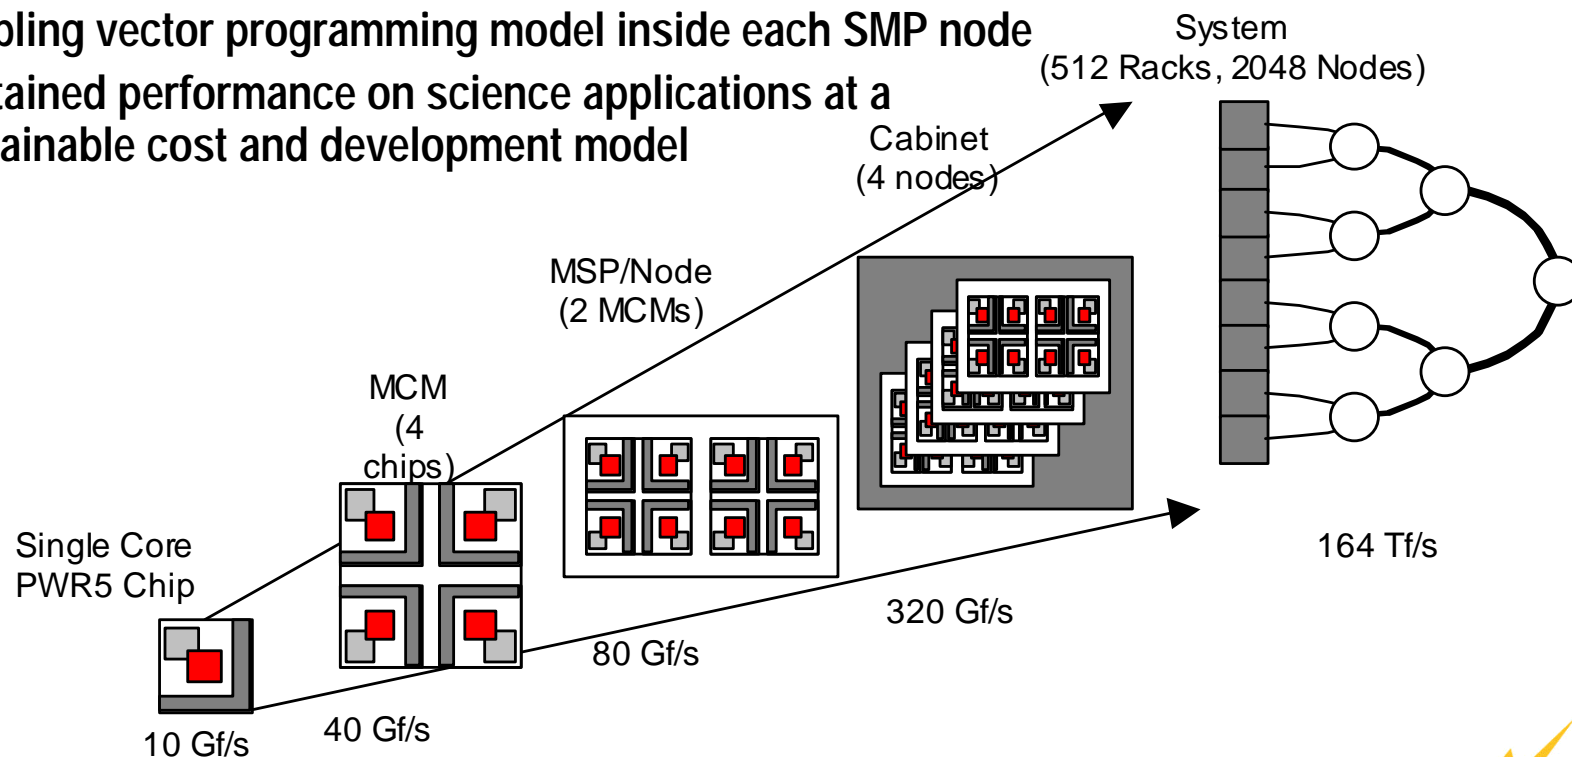




Blue Planet: A Conceptual View



- Increasing memory bandwidth – single core
 - 8 single CPUs are matched with memory address bus limits for full memory bandwidth
- Increasing switch bandwidth – 8-way nodes
- Decreased switch latency while increasing span
- Enabling vector programming model inside each SMP node
- Sustained performance on science applications at a sustainable cost and development model





Progress Already

Even without explicit new funding



- **IBM**
 - **Additional changes Power 5 CPU**
 - E.g. improved memory affinity
 - **Small, more memory intensive node (that actually will cost less than what is proposed in the paper)**
 - **Scaling the switch larger**
 - New ideas for adaptors
 - **Software performance changes**
 - **Significant influence for the Power 6 CPU and chip design**
- **Other vendors and groups engaged**
 - **LLNL**
 - **Vendors**
 - SGI, Intel, AMD, Cray, Others
 - **Modeling**
 - LLNL, NCAR, SDSC, PERC





Scientific Results using NERSC

NERSC Support Efficient Science of Scale

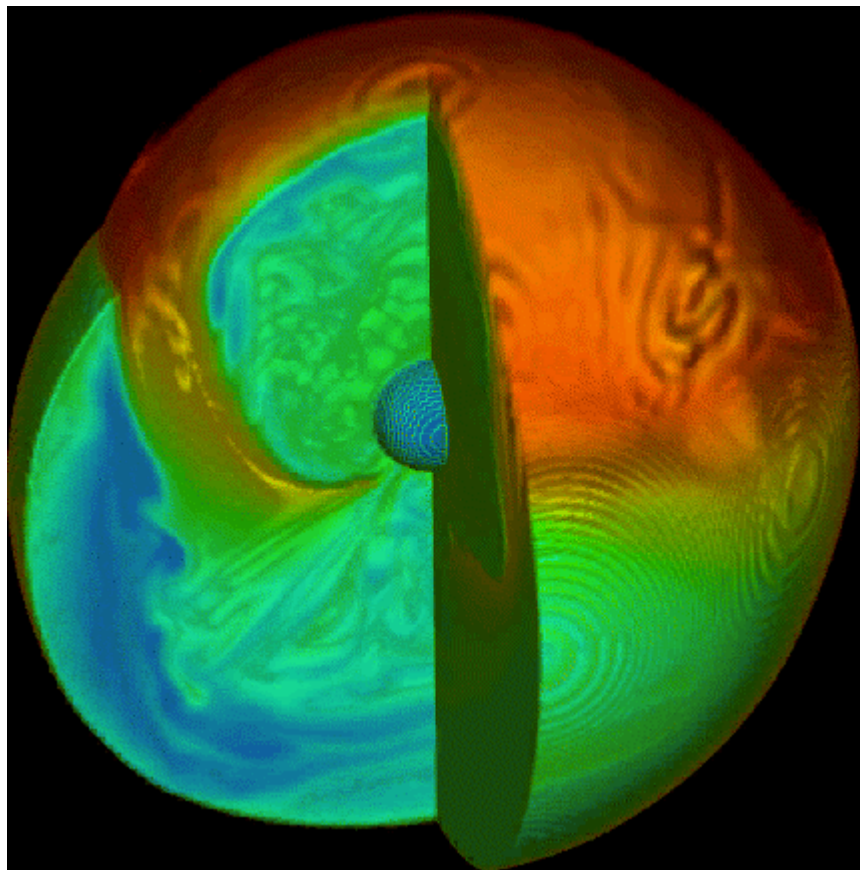
<u>Project</u>	<u>Performance</u> (% of peak)	<u>CPU Count</u>
Terascale Simulations of Supernovae	35%	2048
Accelerator Science and Simulation	25%	4096
Electromagnetic Wave-Plasma Interactions	68%	2048
Quantum Chromodynamics at High Temperature	13%	1024
Cosmic Microwave Background Data Analysis	50%	2048 & 4096

(pre and post processing)

Note – these are comparable to the best documented efficiencies of the science codes on the Earth Simulator, but on different codes of course.



Terascale Simulations of Supernovae



- PI: Tony Mezzacappa, ORNL
- Allocation Category: SciDAC
- Code: neutrino scattering on lattices (OAK3D)
- Kernel: complex linear equations
- Performance: 537 Mflop/s per processor (35% of peak)
- Scalability: 1.1 Tflop/s on 2,048 processors
- Allocation: 565,000 MPP hours; requested and needs 1.52 million



Multi-Teraflops Spin Dynamics Studies of the Magnetic Structure of FeMn and FeMn/Co Interfaces

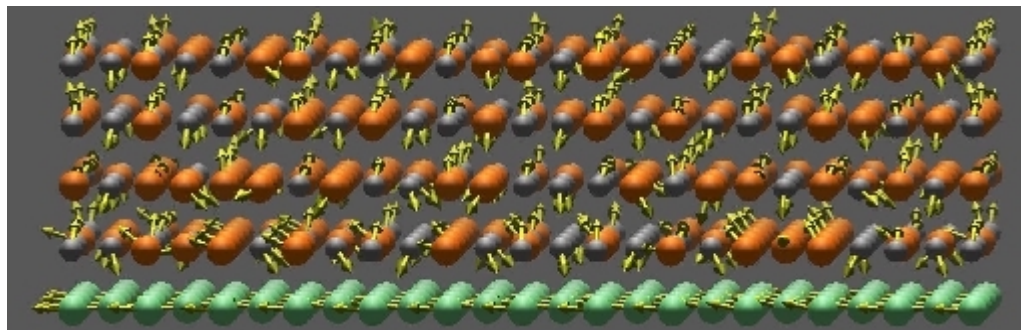
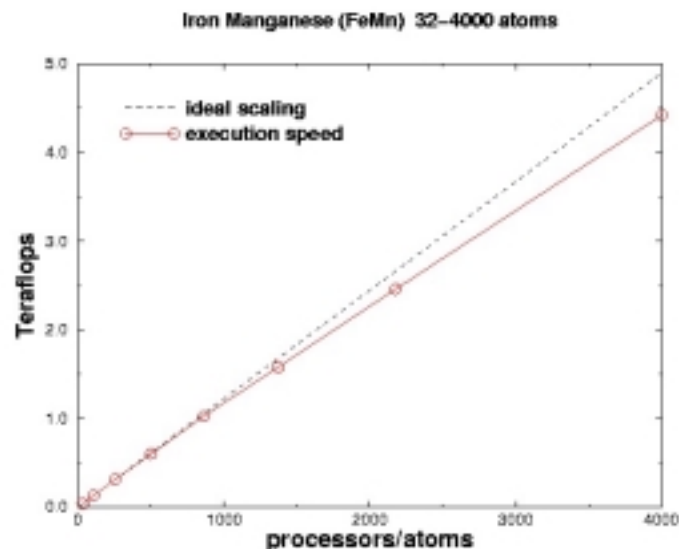


Exchange bias, which involves the use of an antiferromagnetic (AFM) layer such as FeMn to pin the orientation of the magnetic moment of a proximate ferromagnetic (FM) layer such as Co, is of fundamental importance in magnetic multilayer storage and read head devices.

A larger simulation of 4000 atoms of FeMn ran at **4.42 Teraflops** on 250 nodes.

(ORNL, Univ. of Tennessee, LBNL(NERSC) and PSC)

IPDPS03 A. Canning, B. Ujfalussy, T.C. Shulthess, X.-G. Zhang, W.A. Shelton, D.M.C. Nicholson, G.M. Stocks, Y. Wang, T. Dirks



Section of an FeMn/Co (Iron Manganese/Cobalt) interface showing the final configuration of the magnetic moments for five layers at the interface.

Shows a new magnetic structure which is different from the 3Q magnetic structure of pure FeMn.



New Results in Climate Modeling

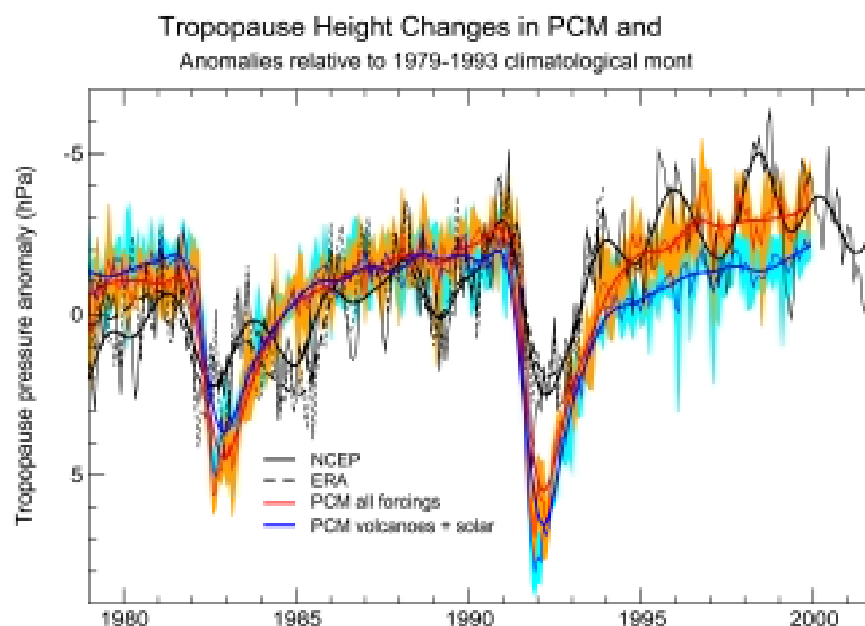
- Recent improvements in hardware have reduced turnaround time for the Parallel Climate Model
- This has enabled an unprecedented ensemble of numerical experiments.
 - Isolate different sources of atmospheric forcing
 - Natural (solar variability & volcanic aerosols)
 - Human (greenhouse gases, sulfate aerosols, ozone)
- Data from these integrations are freely available to the research community.
 - By far the largest and most complete climate model dataset
 - www.nersc.gov/~mwehner/gcm_data



Investigating Atmospheric Structure Changes with PCM



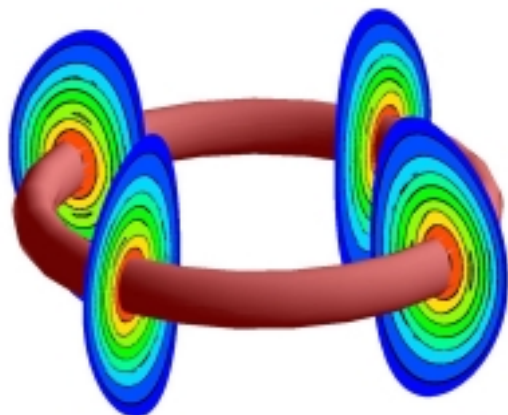
- The tropopause is that height demarking the troposphere and the stratosphere.
 - Below the tropopause, the temperature cools with altitude.
 - Above the tropopause, the temperature warms with altitude.
- A diagnostic that is robust to El Nino but sensitive to volcanoes.
- An indicator of the total atmospheric heat content
- Changes in natural forcings alone (blue) fail to simulate this feature of the atmosphere, but natural + anthropogenic changes (orange) do



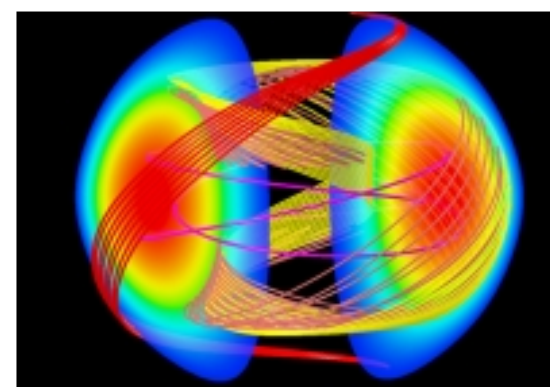
Santer et al. Figure 1



SCIDAC Collaboration Speeds Up Fusion Code By Factor of 10



- NIMROD is a parallel fusion plasma modeling code using fluid-based nonlinear macroscopic electromagnetic dynamics.
- Joint work between CEMM and TOPS led to an improvement in NIMROD execution time by a factor of 5-10 on the NERSC IBM SP.
- This would be the equivalent of 3-5 years progress in computing hardware.
- Parallel SuperLU, developed at LBNL, has been incorporated into NIMROD as an alternative linear solver.
 - Physical fields are updated separately in all but the last time advances, allowing the use of direct solvers. SuperLU is >100x and 64x faster on 1 and 9 processors, respectively.
 - A much larger linear system must be solved using the conjugate gradient method in the last time-advance. SuperLU is used to factor a preconditioning matrix resulting in a 10-fold improvement in speed.





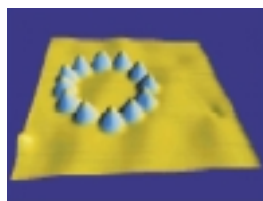
Projects in Computational Research Division

--

Bringing all resources together

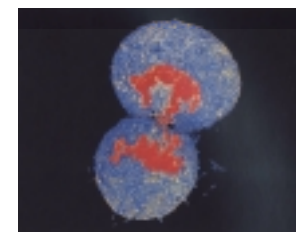


Computational Science Mission

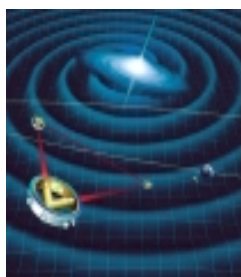


**nano
systems**

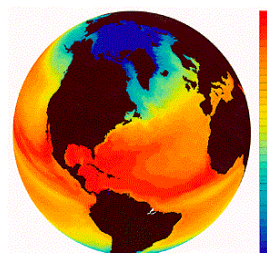
**biological
systems**



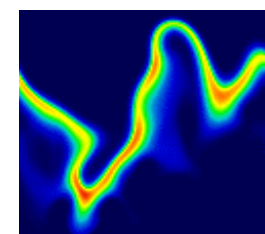
**The Computational Research
Division is engaged in
computational science
collaborations, creating tools and
techniques that will enable
computational modeling and
simulation, and lead to new
understanding in areas such as**



**astrophysics
simulation**



global climate



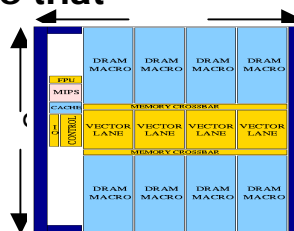
**combustion
processes**



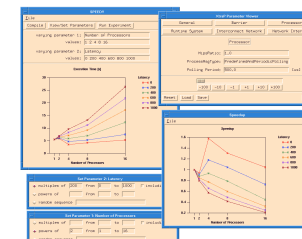
Computer Science and Applied Mathematics Mission

The Computational Research Division is engaged in basic and applied research addressing the following questions of fundamental importance to enabling progress in our ability to use computing and networking technology

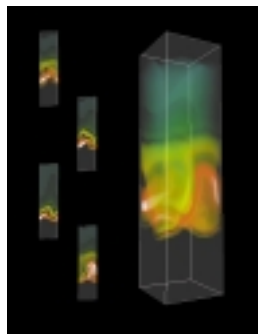
- develop computer architectures that most suitable for scientific applications and measure their effectiveness for science



- research in algorithms and development of software tools for these new architectures



- research in algorithms, and development of software tools, and technology in data management, analysis, and visualization

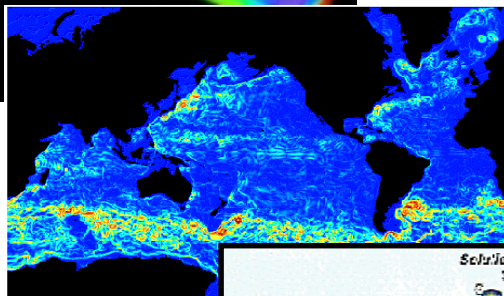
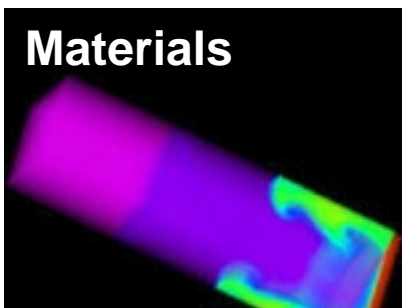


- research in networking and distributed computing, and development of grid middleware and collaboration technologies

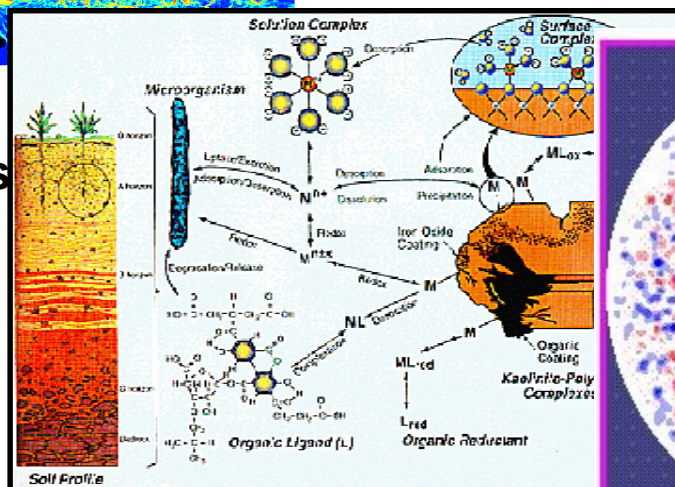


Scientific Discovery Through Advanced Computing

Materials

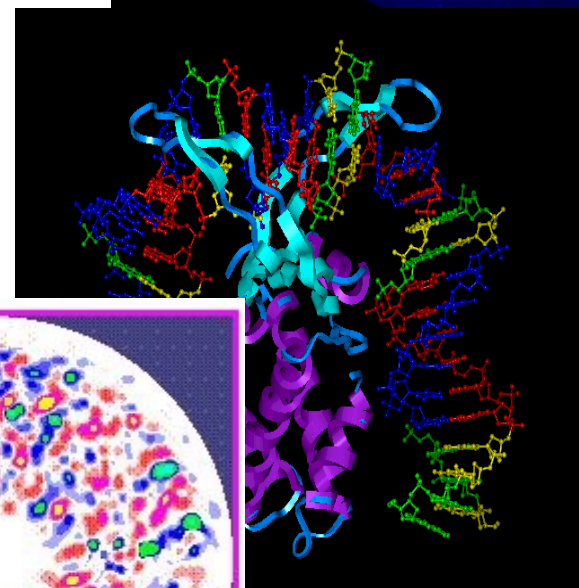
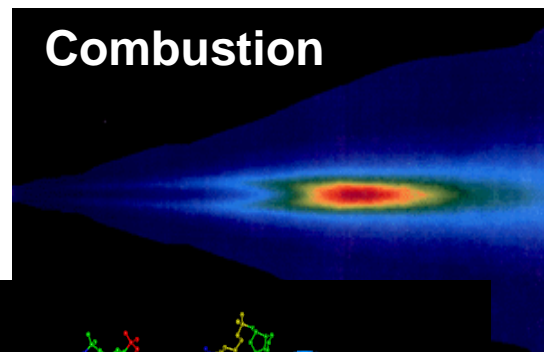


Global Systems

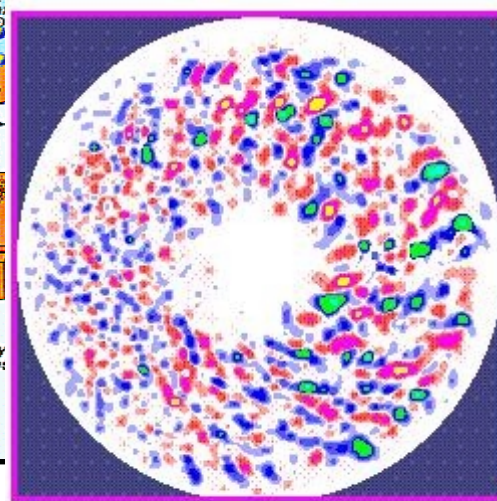


Subsurface Transport

Combustion



Health Effects, Bioremediation



Fusion Energy

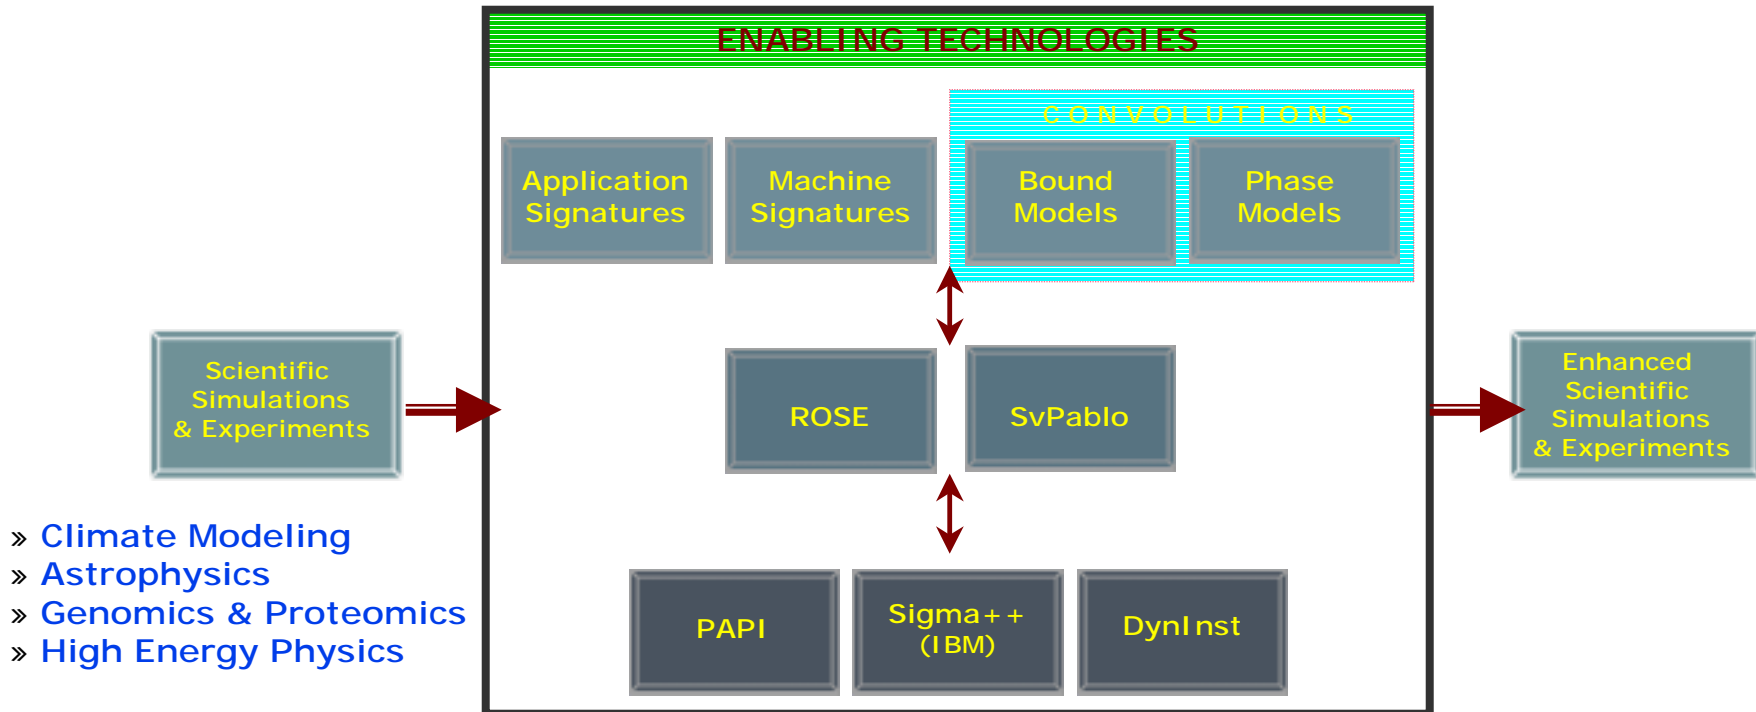
Performance Evaluation Research ISIC

Developing a *science* for understanding performance of scientific applications on high-end computer systems, and *engineering* strategies for improving performance on these systems.

GOALS

Optimize and Simplify:

- Profiling of real applications
- Measurement of machine capabilities
- Performance prediction
- Performance monitoring
- Informed tuning



Participants: ANL, LBNL, LLNL, ORNL, UC San Diego, U.I.-UC, UTenn



Performance Evaluation



- Signed MOU with Earth Simulator Center about joint evaluation of vector architectures for scientific applications
- Collaboration with ORNL about Cray X1 evaluation
- Alternative architecture evaluation project:
 - VIRAM (Patterson, Yelick), DIVA(Hall), Streams (Dally)

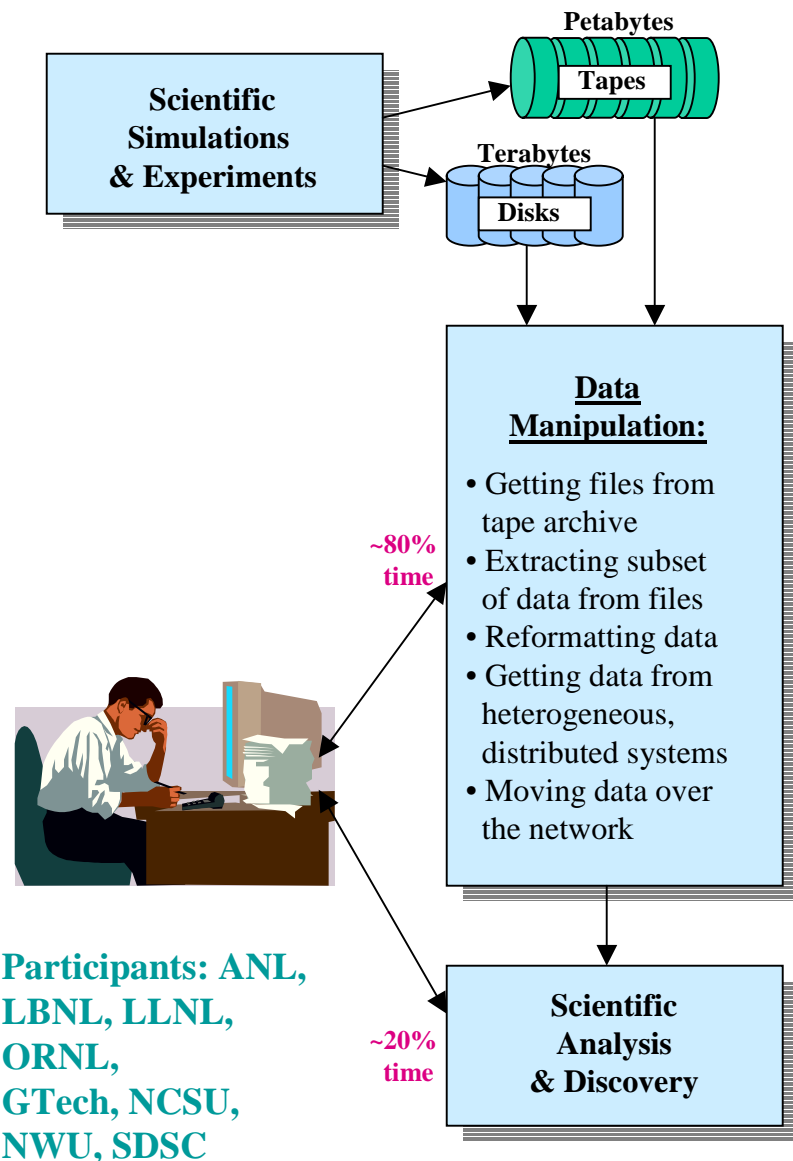


Scientific Data Management ISIC

Goals

Optimize and simplify:

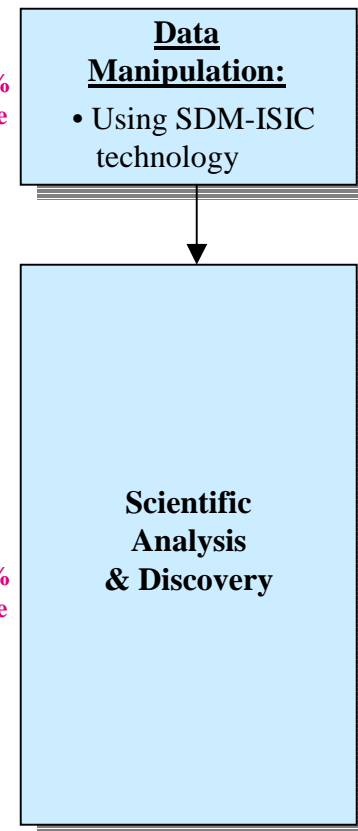
- Access to very large data sets
- Access to distributed data
- Access of heterogeneous data
- Data mining of very large data sets



SDM-ISIC Technology

- Optimizing shared access from mass storage systems
- Metadata and knowledge-based federations
- API for Grid I/O
- High-dimensional cluster analysis
- High-dimensional indexing
- Adaptive file caching
- Agents

~20% time



The DOE Science Grid: A New Type of Infrastructure

Early Example:
Access Grid –
Multi-site, high-end
collaboration

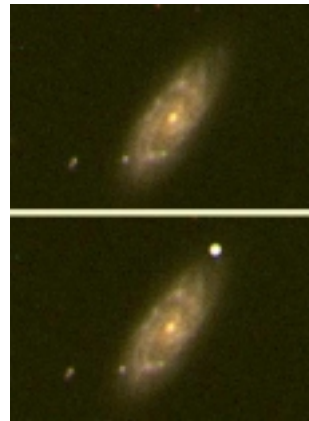


Lab, University, and Industrial partners
collaborate using unique instruments, large
amounts of data, and computing facilities at
multiple sites.

Images &
Spectral Data



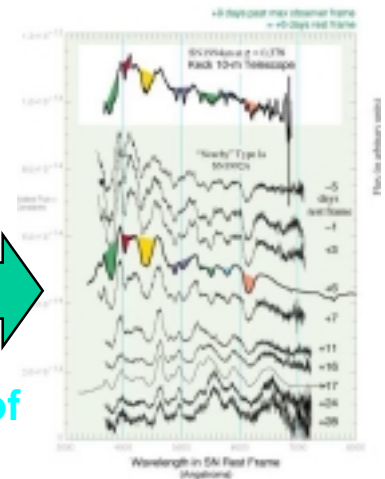
Telescope
retargeting



Supernova
identification



Simulation of
Spectra



Example: The Supernova Factory (SNAP)

Participants: LBNL, ANL, ORNL, PNNL (and NERSC)

Simulation Science by 2012: Cosmic Simulator



Science driven vision of a computational framework in 2010.

The Cosmic Simulator is the concept of providing an integrated framework in which component simulations can be linked together to provide a coherent, end-to-end, history of the Cosmos.



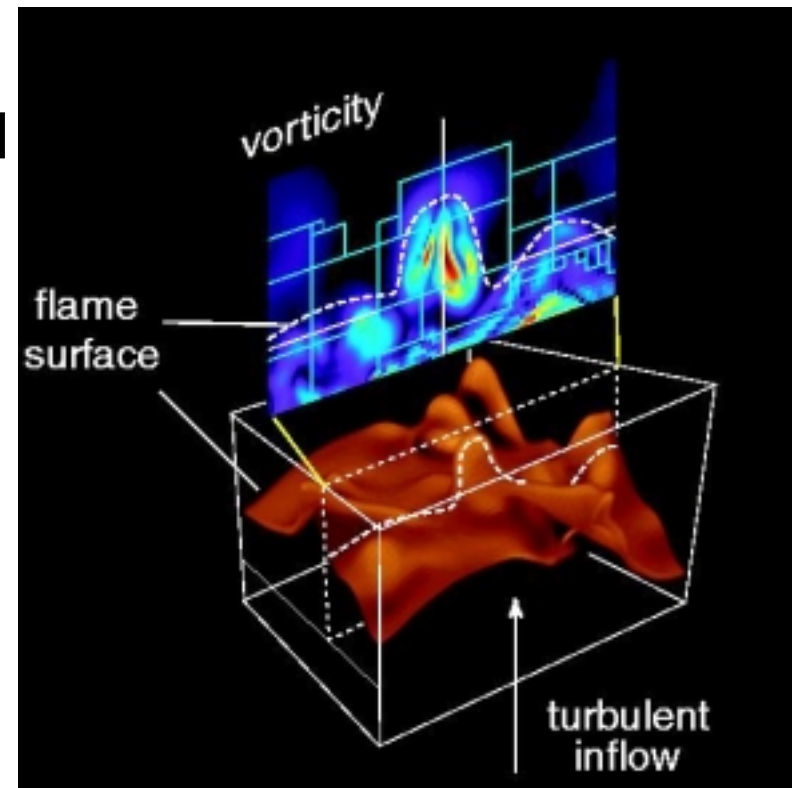
South Pole



Applied Partial Differential Equations ISIC

Developing a new algorithmic and software framework for solving partial differential equations in core mission areas, such as accelerator physics, magnetic fusion, and combustion.

- New algorithmic capabilities with high-performance implementations on high-end computers:
 - Adaptive mesh refinement
 - Cartesian grid embedded boundary methods for complex geometries
 - Fast adaptive particle methods
- Close collaboration with applications scientists
- Common mathematical and software framework for multiple applications



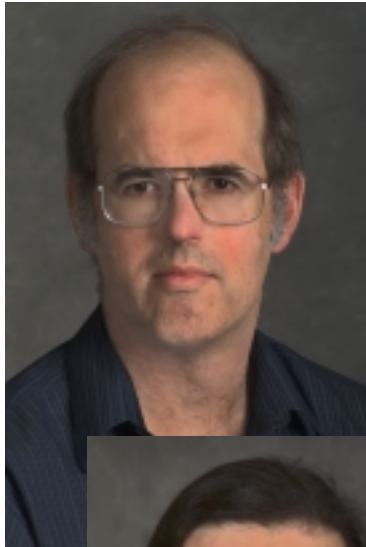
Participants: LBNL (J. Bell, P. Colella), LLNL , Courant Institute, Univ. of Washington, Univ. of North Carolina, UC Davis, Univ. of Wisconsin.



2003 SIAM/ACM Prize in Computational Science and Engineering

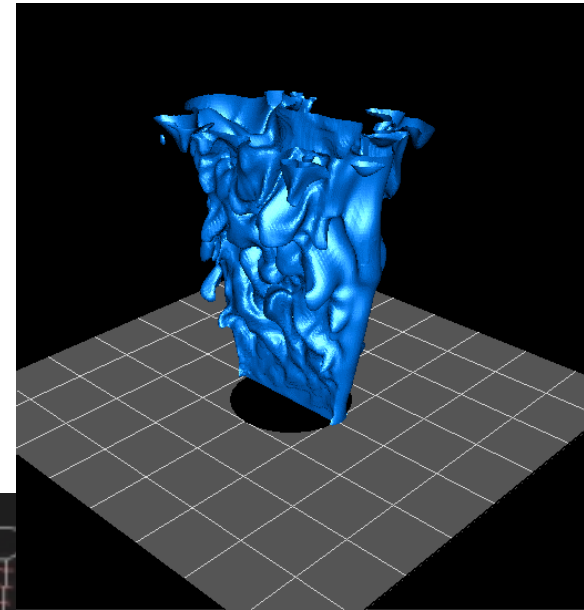


Awarded by the Society for Industrial and Applied Mathematics (SIAM) and the Association for Computing Machinery (ACM)

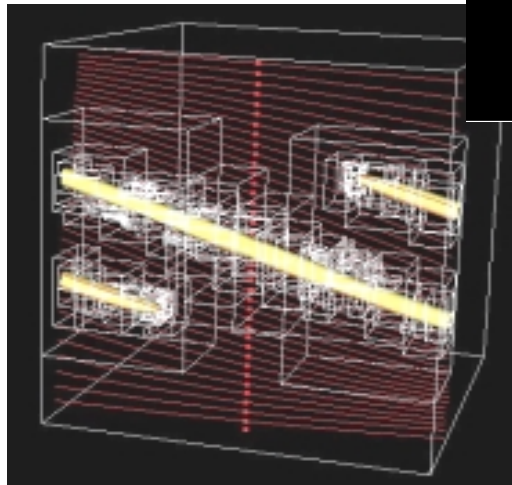


John
Bell

Chemistry and
turbulence in
methane
combustion



Phillip
Colella

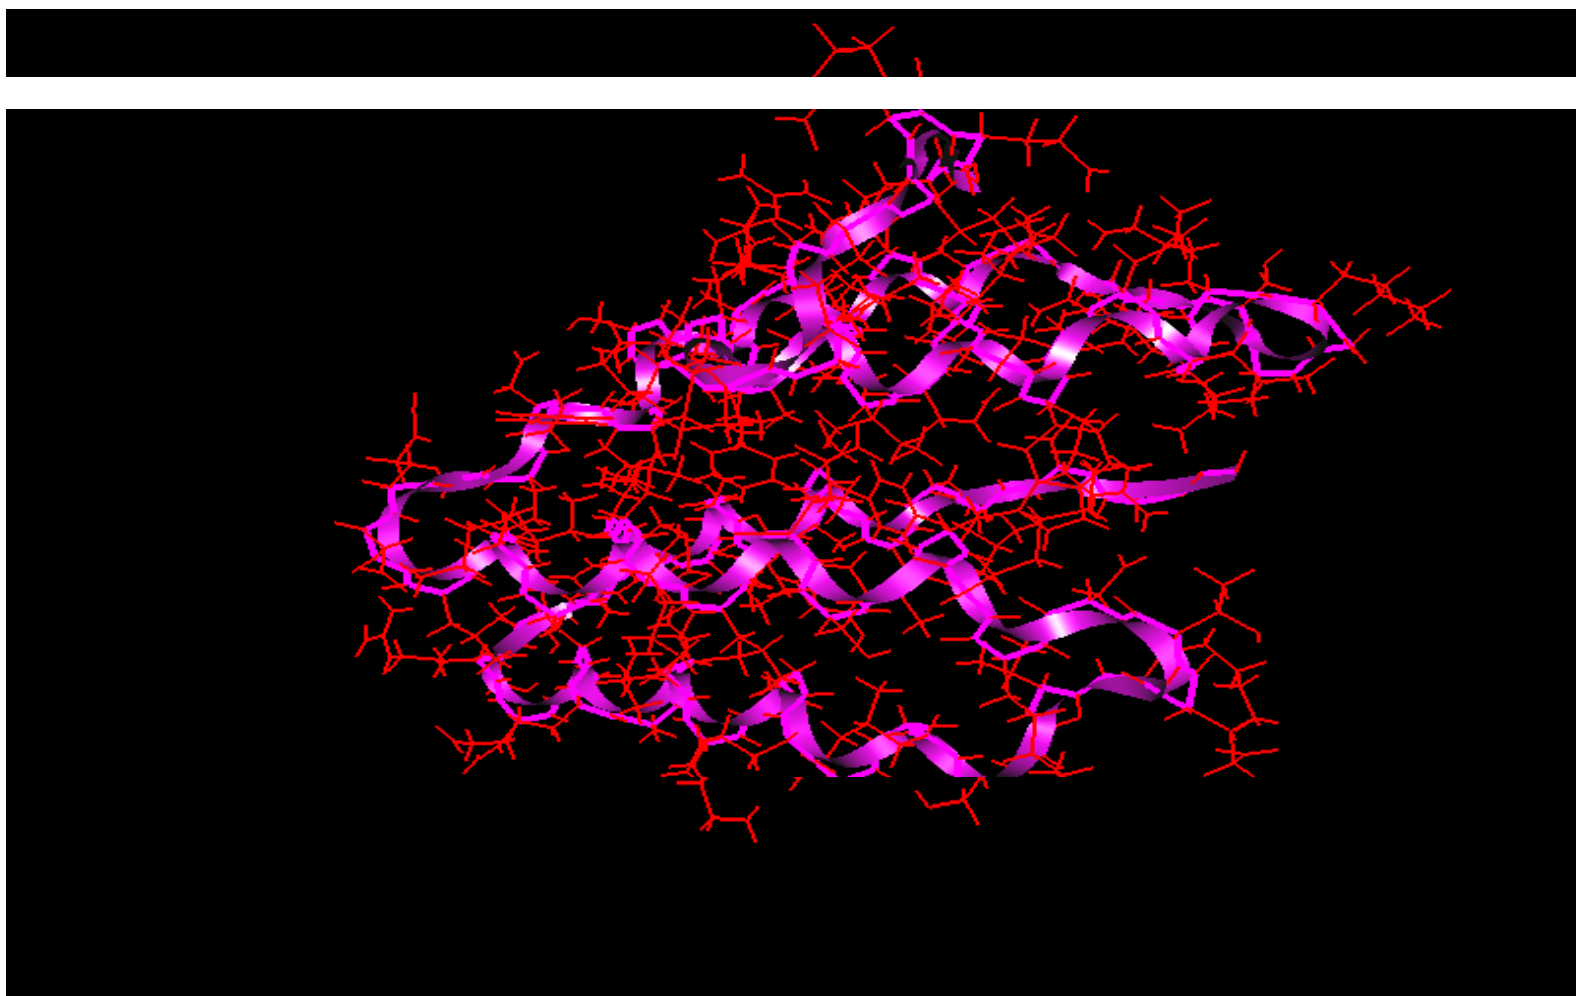


Three-dimensional
adaptive mesh
calculation of pellet
injection into a plasma





Computational Steering of Protein Folding



Source: Silvia Crivelli et al. , CASP 2002

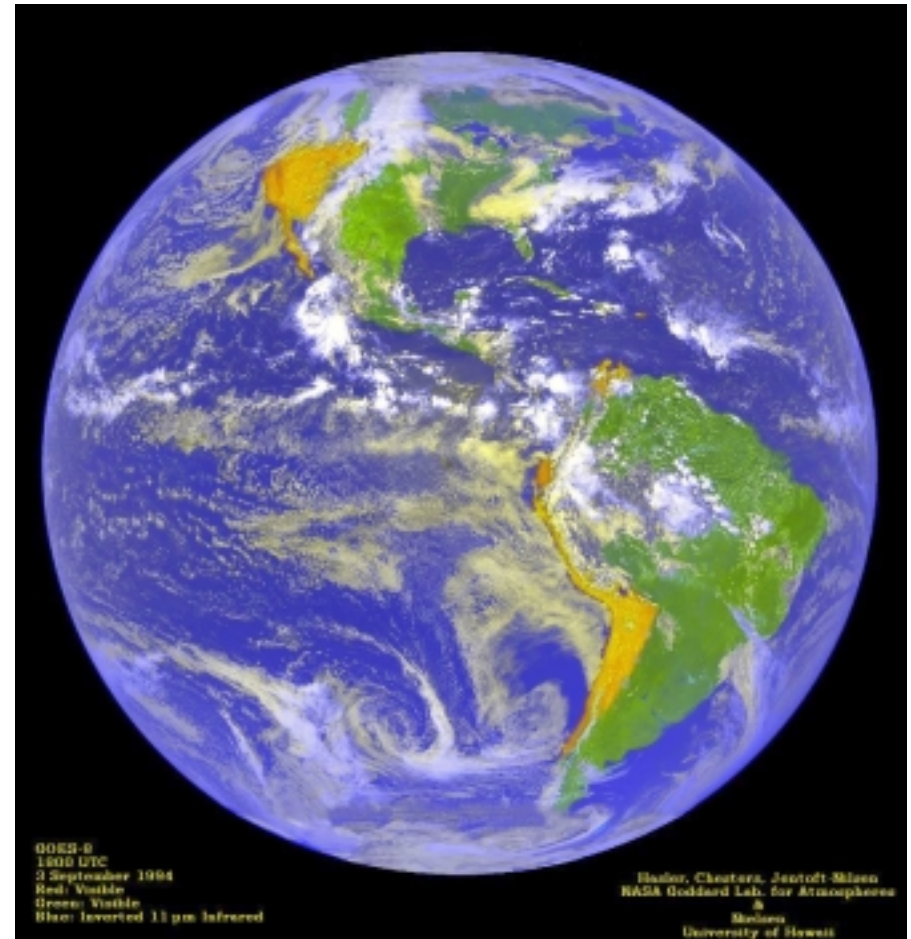




JGI: Sequencing the Biosphere



- Less than 1% of microbes are culturable
- Many unculturables live in diverse interdependent consortia
- Aim:
 - Recover genome-scale sequences and reveal metabolic capabilities
 - Understand action of microbes at molecular level
 - Determine structure of microbial populations



Source: Eddy Rubin, JGI

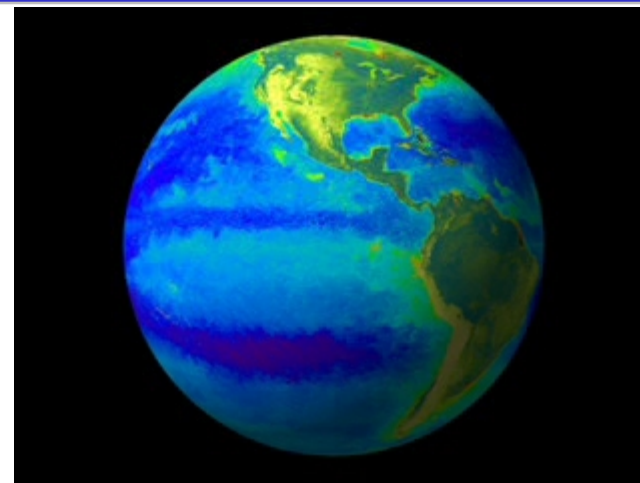




Collaboration with JGI in Data Management



- Genomic diagnosis of environment
- JGI as a National User Facility
 - 40% for DOE missions
 - 60% for scientific community
 - Department of Agriculture: crop pathogens
 - Environmental Protection Agency: environmental indicator species



Fathead
Minnow



Sudden Oak
Death

- FY2004: data pipeline from JGI to NERSC HPSS

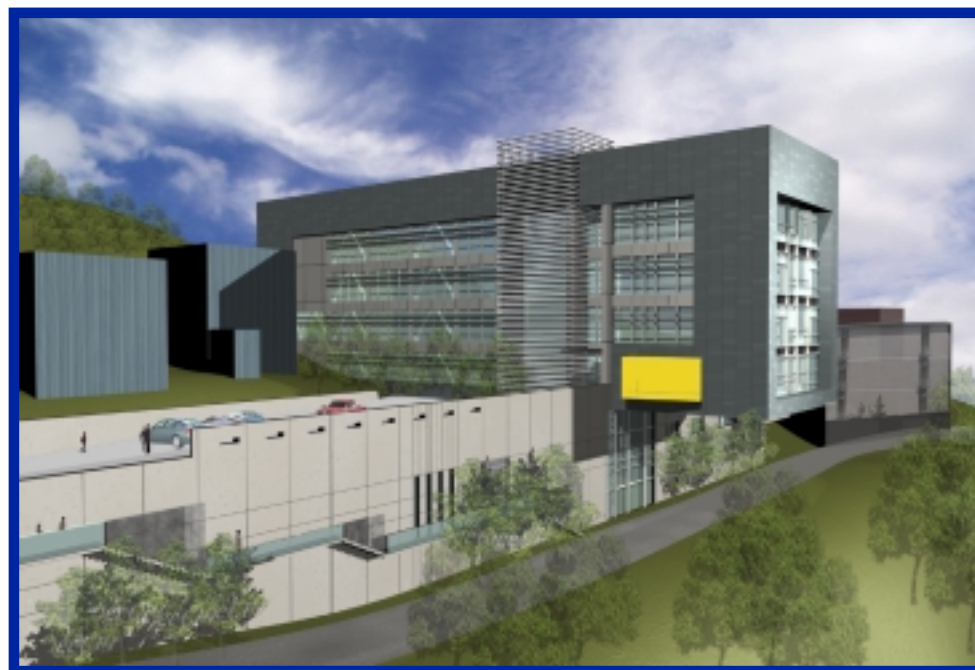




Molecular Foundry Status

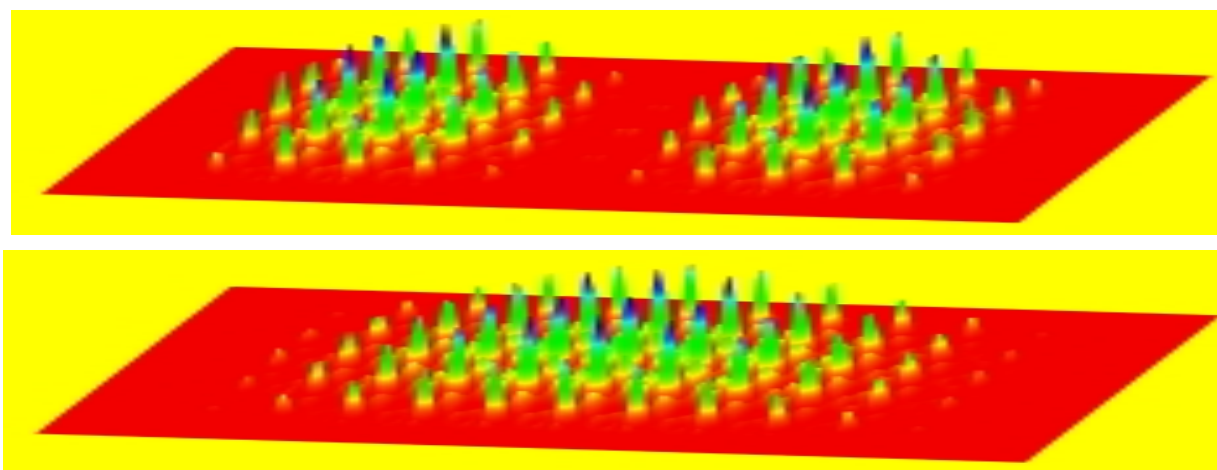
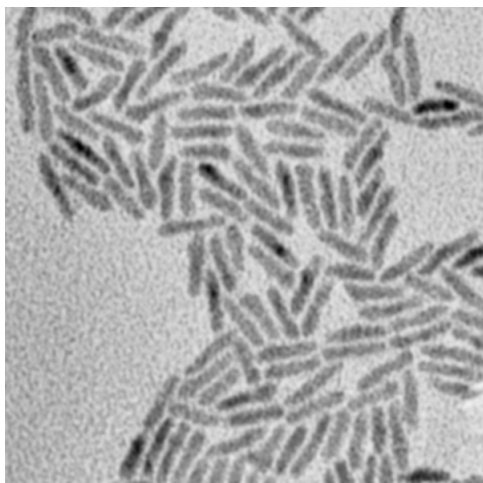


- Design approved by Regents, April 2003
- Successful independent external review completed, June 2003
- Budget:
 - \$6.8M for FY 2003
 - Total project cost: \$85M





Collaboration with Molecular Foundry in Computational Nanosceince



CdSe quantum rods

The electron wavefunctions of a quantum rods

- The electronic structures and optical properties change with the shape of the quantum rods.
- The thousand atom quantum rods can be calculated using the planewave pseudopotential method and the NERSC supercomputers.
- Programs exist at NERSC to calculate such nanosystems and compare with experimental electronic and optical results.



The Future



More resources:

No limits to growth in demand for supercomputer resources seen

Better integration:

Computational science and engineering will become recognized as discipline

Next level simulation science:

Large scale simulation environments will emerge that allow computer simulation at unprecedented scale





Summary



- NERSC is one of the largest open, unclassified supercomputer centers world wide, and is one of the largest “data centers” in science
- CRD has unique capabilities in
 - Distributed computing and grid middleware
 - Data management
 - Cybersecurity
 - Imaging and Visualization
 - Scientific Computing and Numerical Methods
 - Applied Mathematics
 - Applied Computer Science Research
- Both leverage DOE/SC investment through
 - Collaborations in a DOE/SC multipurpose lab
 - International Collaborations
 - University Collaborations

